Fermi National Accelerator Laboratory

FERMILAB-Conf-84/27
2300.000

THE FUTURE OF COMPUTING TECHNOLOGY IN PHYSICS - THE POTENTIALS AND PITFALLS*

A. E. Brenner

February 1984

*Invited talk, presented at the 1984 Annual Joint APS/AAPT Meeting, San Antonio, Texas, February 2, 1984.

THE FUTURE OF COMPUTING TECHNOLOGY IN PHYSICS- THE POTENTIALS AND PITFALLS*

A. E. Brenner, Fermilab

A. Introduction

No one will challenge that the technological era in which we now live is the age of the computer. It may be less clear that the effects on our society have been enormous heretofore. I venture to guess that historians in the 21st century will assess that the transformations wrought by the innovations of this age were at least as important and dramatic as those brought about by the Industrial Revolution. As with any profound change in society, driven by motives to improve that society, the positive transformations are inevitably accompanied by a set of secondary effects which at least some segment of society recognize as undesirable. Even for those effects of the Industrial Revolution that all would agree are negative, e.g., Love Canal, acid rain, the pollution of our air and water, reasonable people will disagree as to how to control the level; whether they're inevitable or not; and what additional price society should be willing to pay to modify the effects. Similar very important questions arise with regard to changes wrought by the age of the computer.

*Invited talk, presented at the 1984 Annual Joint APS/AAPT Meeting, San Antonio, Texas, February 2, 1984.

Since we are participating in this transformation of our society, we may affect, indeed have a responsibility to influence, the directions in which the transformation takes place. But as with the negative effects of the Industrial Revolution, we will certainly not all see and interpret the secondary and higher order effects in the same way. Nevertheless, one goal of this paper is to describe some of the pitfalls in this transformation as I perceive them, as I also summarize the fantastic opportunities of this age.

B. Background

As with all other human endeavors, it is difficult to pinpoint the beginning of the computer age. However, let me choose the early nineteenth century, the time of Charles Babbage, as the beginning of the current era. Of course, the early seeds go back to the ancients who built devices, such as Stonehenge, to aid them in calculating the mathematical problems of importance to them. After a long childhood, one might say that the adolescent spurt started in the mid 1940's. During the early days of this period, once the stored program idea had been formulated by John von Neumann and his colleagues, the primary thrust was to learn to develop and organize the technology to develop machines capable of removing the drudgery part of the scientific communities' requirements to compute, or "number crunch," and the more general requirement to acquire and manipulate large volumes of scientific or business data.

Today with the maturation of silicon-based technology, with over 25 years experience in learning how to mass produce and make this technology work effectively, it is now possible to place on a single chip, smaller than a fingernail, more computing power than filled a room in the days of the first commercial modern scientific computers, the Univac I and the IBM 704. Furthermore, the cost of this computing power has gone from several million dollars to several tens of dollars.

One measure of growth during this early period is the raw computational power of a single computing engine or Central Processing Unit (CPU) produced by industry. Although this measure is now becoming less meaningful, it has been a good measure during most of this adolescent period and Fig. 1 traces that history. The units are millions of instructions per second (MIPS) or in "number crunching"-type problems, since one is primarily limited by the rate of arithmetic computations, one uses a measure of millions of floating point operations per second (MFLOPS). Note that by this measure the rate of growth is slowing down now relative to where we were earlier during this period. Nevertheless, there still is enormous growth in front of us which this figure does not properly represent. The computers represented here are all classical von Newmann architecture or serial execution machines. New technologies and especially new architectures will once again increase the rate of growth during the next decade.

Through the first three decades of this adolescent period, the motivations and the efforts were focused on making more powerful CPU's. Although the development of larger and more powerful "number crunchers," the supercomputers, continues today, and is very important to the scientific communities' needs and more generally to societies' needs, a new facet is emerging as the aspect of computing development receiving the most attention. Furthermore, it is probably the most important aspect of this technological revolution, both in its positive and in its negative attributes. This is the development of workbench or, more generally, productivity enhancing tools. This development has more or less been coincident with the advent of the microprocessor and the success of very large scale integration (VLSI) techniques. Whereas in the early days of the computer era the focus was to obtain computing cycles, with which to compute or move data, the major focus today is in the human interface to make whatever computing capabilities are available "user friendly." This is only possible now because that powerful inexpensive chip is readily available. As with the earlier number crunching and data moving developments, this new aspect of computer development is having and will have an enormous effect on the way in which we do our physics and, more generally, on society.

The productivity enhancing tools themselves do require enormous amounts of computing capability unseen by the user. Although the type of functions performed are typically more decision oriented and less arithmetic, nevertheless large numbers of operations are required to make an interface with the user friendly. For example, the computer langauge BASIC is popular for good reason. It is simple to use and for most simple computations responds as quickly as any user would like, even when executed on the least expensive personal computer. It and other similar "interpretive languages" used interactively have excellent user friendly attributes. However, run a major computation bound program written in BASIC and one would find one's throughput is degraded by many orders of magnitude. The reason for this is that interactively using interpretive mode programs, while making life simple for the user, does require large numbers of additional functional cycles above and beyond that required for the primary calculation.

In addition to the advent of the small, powerful processor on a chip, another important maturation has been occurring and is now beginning to have an important effect. That is the understanding of the software problem. In the early days of this age, it was the hardware which required attention, development and understanding. Software was not thought to be a problem. The first major scientific computers delivered, the Univac 1 and IBM 704, were delivered to their first customers with absolutely no software, not even the software required to run peripheral devices such as card readers and printers. The importance of software made its first major impact in the mid 60's when IBM, after a well known internal corporate struggle,(1) chose to introduce a new line of equipment, the System 360, which is basically the system IBM manufactures today. Given their enormous experience and competence, IBM had fully understood[2] the cost of development of the hardware. However, the software problem had been underestimated in cost by more than one order of magnitude and if the accounting is properly done, probably by two orders of magnitude. A company less than IBM certainly would not have made it through that period, as they expended enormous resources to reinvent a large number of tools which had been developed both by IBM and its customers during the period of the earlier generation of IBM computers.

Software has come a long way since then. However, it still is an unusual commodity with attributes not very different from our spoken and written language. In the sense that a baby can speak the same language as the poet, there is a similar difference in the power, the content, the efficacy and the elegance in the programs written by an amateur and a professional. But it is the availability of software and our better understanding now as to how to manage and use it that has made it possible to develop effective work enhancing tools. It is also helping in the development of new architecture supercomputers.

Finally, one other technological development is relevant to our picture. That is, the integration in a coherent and natural way of computers and the communications tools that society has been building for the last 100 years. This integration[3] will place the computer in a yet more important position in our lives and work. Although ideas of this integration were already articulated in the late 50's by McCarthy and Corbato[4] in their arguments for what became MIT´s project MAC, the details are turning out to be rather different. Their concept at that time was to share the power of computers across the nation in the same way in which generating stations share the power in our national electrical grid. That has not worked, nor will it work. But it is the sharing of information, the other facet for which the computer has become so important, where the communications networks play an exceedingly important role.

C. The Tools

The advent of the current generation of supercomputer, the Crays and the high-level Cyber computers, which for some specialized problems compute in the hundred millions of floating point operations per second, is changing the character and style in which many engineering and physics problems are solved. With the current capabilities of these supercomputers, any number of heretofore intractable problems become soluble. Thus, it is possible today to simulate rather extensively activities that previously could only be modeled using analog techniques in the laboratory. A trivial example is that it is now possible and necessary to simulate very large scale integrated circuits to produce successful large scale circuits with finite effort. For aircraft design, it is far cheaper to buy a supercomputer than to build another wind tunnel and to simulate a particular air foil before the prototype is built. The efforts today to design a superconducting supercollider accelerator, the SSC, 40 kilometers in diameter, costing $2-3 billion is only feasible because the more complex components and the system in its entirety will be simulated before one actually commits to their construction.

The mini-computer has broadened its role over the years. It now extends into the realm of the larger general-purpose processors (the mega-mini) and also down to the personal computer (the micro-computer). These tools have been particularly useful in the past to the experimental community for data acquisition and for the control of equipment and experiments. With advancing technology, these computers have grown in power and are even more useful in this role to the experimental researcher. However, their power has increased to such an extent that they are now also beginning to be of interest to the theorists.

Of course, with the price so low at the personal computer end of the spectrum, we find schools in America, from the local grade school through our great universities, acquiring these in growing numbers. Indeed, there is a great give-away race in progress as the various manufacturers of these units vie for the loyalties of our young people___who will make or break the future of most of these competing firms.

The recent advent of the scientific workstation, making available at a scientist's desk a working tool capable of major computational power to manipulate algebraic quantities symbolically, to perform a demanding numerical calculation, to manipulate a data set and to graphically display the results will improve both the productivity and the effectiveness of each researcher. Furthermore, the cost of these units is low enough that the smallest of research projects should be able to take advantage of such tools.

The same technology, which is mass producing small sized, inexpensive and highly capable logic systems, also gives rise to special purpose processors, very powerful and capable of doing many complex functions which in the past were either hardwired or not done at all. Now it is possible to embed large numbers of powerful computing engines with large memories to make computational decisions concerning complex physical events in times of less than one millisecond. This has made possible complex experimental set-ups that were impractical heretofore.

D. Leading Edge Examples—The Potentials

The physics community was one of the very first major users of computers in this computer era. Although no longer the pre-eminent user, the community continues to push the technology at the leading edge. This is certainly true in the area of both number crunching and data acquisition. Although we may not be at the very leading edge relative to the work enhancing tools,[5] we are certainly in the first rank.

In many of our enterprises, there is a tight coupling between the data acquisition and number crunching uses of computing technology. For data acquisition, we have used and continue to use mini-computer processors. For the data analysis or number crunching phase of our research, depending upon the details of the problem, we properly use dedicated mega-mini computers, powerful general-purpose computers or the supercomputers of the day. Each field of research has a different style in this arena, changing with time, but mostly dictated by the need and the cost and availability of the resources. The fantastic rate of growth of our computer technology, pushed in part by our needs, but basically paced by the engineering and manufacturing capabilities of industry, has almost kept up with our needs. With these tools, we have been able to take on investigations that were not at all possible before these tools became available. To best illustrate the potential impact of this technology on physics, a number of leading edge examples of instruments just beginning to work or scheduled to be finished in the near-term future are given here:

(1) The Very Large Array

The Very Large Array (VLA) radio telescope situated in the Plains of San Augustin near Socorro, New Mexico, see Fig. 2, consisting of 27 dishes, each 25 meters in diameter located along the legs of a Y, each arm of which is 21 kilometers long, uses a number of mini-computers for control and data acquisition. During some measurements, the continuous data acquisition rate is 300K bits per second. The data is processed by array processors and other special-purpose processors developed for this function. Currently, the data is analyzed on a large mid-sized computer. This instrument is already severely limited in its capabilities because of its inability to analyze all the data collected. Plans are currently being developed to acquire a supercomputer for this analysis role.

(2) Controlled Fusion

In the controlled fusion program, two examples are the Princeton Tokamak Fusion Test Reactor (TFTR), Fig. 3, and the Livermore double Mirror Fusion Test Facility (MFTF-B), Fig. 4. In both cases, a number of mini-computers are utilized for control and data acquisition. The data acquisition rates in this case are quantized on shot times and rise to over 300 Mbits/shot; the frequency of shots varies from several minutes to several hours depending upon experimental details. For these fusion experiments, the analysis at the first level is performed on mid-sized computers. Frequently, however, the analysis requires the use of supercomputers. The Magnetic Fusion Energy Computing Center (MFECC) located at the Lawrence Livermore Laboratory, with its two Cray 1 supercomputers, was established to support such activity across the country. As an integral part of the fusion research program, the theoretical work done in association with these experiments are some of the most demanding number crunching research programs in our field.

(3) The Space Telescope

The space telescope shown in Fig. 5 with its 2.4 meter fused silica mirror is currently scheduled for orbit in 1985. This instrument also demands an intimate integration of complex electronic components, broadband communications, and data acquisition and analysis computers. The data stream for scientific data is semi-continuous and exceeds 1 Mbit/sec.

(4) DUMAND

DUMAND is a deep (4.7Km) underwater detector, 250x250x500m^3, consisting of an array of 756 16" diameter photomultipliers shown in Figs. 6 and 7. The scientific aim of the project is to study high energy neutrino astrophysics, high energy neutrino particle physics, cosmic ray physics, and ocean and earth science. Data from each string of 24 detectors is carried over fiber optic cables 25km to shore continuously at 44MHz. Each detector has a small processor in the Benthos sphere vessel holding the detectors. There are special-purpose processors at the bottom of the ocean to format the data and control the detector. On shore, there is an array of special-purpose processors to analyze the data.

(5) High Energy Physics

As a final example of experimental science, I would like to focus in a little more detail on the use of computers in high energy physics. Here also there is a fundamental need for computer technology integrated at all levels. Numbers of powerful mini-computers control and run our accelerators. An overview of Fermilab showing the 2 kilometer diameter main accelerator ring is shown in Fig. 8. A schematic of the Fermilab accelerator control system is shown in Fig. 9. For the experiments at this and other accelerator laboratories, dedicated large mini-computers are utilized for each experiment. Table 1 shows the approximate data rates as a function of time for high energy physics experiments.

New large detectors are being built at CERN for the new 8.5 kilometer diameter electron-positron colliding beam accelerator (LEP) now under construction, and for the collider facility being completed at Fermilab. The Fermilab Collider Detector Facility (CDF) is shown in Fig. 10. As an example of the complexity of some of the data from these detectors, a typical event from the UA1 detector at CERN is shown in Fig. 11. This generation of experiments typically acquires data at the rate of 10^5 bits/event with event rates sometimes as high as 100 per second. The actual stored data rate is invariably limited by the rate at which current storage technology allows, typically 6250 bpi tapes running at their maximum speeds. To get to those rates, an enormous amount of processing must be done on-line to reject those events which are less interesting or less relevent for the particular physics under study. To do this requires embedding in the data acquisition electronics large numbers of either special processors or special-purpose electronics to make quite complex decisions in very short times. Thus, in addition to 5 nanosecond circuit decisions made by hardwired elements, there are any number of decisions made in many of these experiments within 1 to 20 microseconds. Typically these require a number of arithmetic calculations and are frequently done either by fast special-purpose processors tuned to the particular needs or by more general-purpose lookup procedures in memories where the solutions for all sets of parameters have been prestored. It is here where the advent of quite inexpensive powerful processors and the technology that has made that possible, also makes it feasible for these experiments to be contemplated and mounted. The planned data acquisition system for the Fermilab CDF detector is shown in Fig. 12 as an example of how these experiments utilize computing technology for data acquisition.

Having written all of this data on tape during the data-taking phases of these experiments, the analysis and number crunching requirements are formidable. These new large experiments, such as the UA1 and UA2 experiments at the CERN SPS, have formidable computing requirements. These two CERN experiments which discovered the W and Z^0 last year have only scratched the surface in the analysis of the data that is already in hand. Available computers at CERN and in Western Europe, more generally, are not at all adequate to the chore. Computers with new architectural designs are needed to handle the volume of computing required.

One solution appropriate to this particular problem is to use numbers of relatively inexpensive processors banked together to operate asynchronously. Some of these have been built, originally at SLAC, emulating the IBM architecture. Quite a number of these have been built now and they are in use in a number of high energy physics laboratories. They have been used at CERN, both on-line in the UA1 detector and off-line as an adjunct to the CERN central computing facilities. This approach is very cost effective. The price one pays, however, is that the network organizational problems are non-trivial and are not yet adequately solved.

(6) Theoretical and Other Computational Problems

In general, theoretical problems solvable by direct computation or by simulation are beginning to become interesting to a much broader range of theoreticians. The more powerful computers have made it possible to do practical computations using lattice techniques to understand fundamental questions in physics from critical phenomena to gauge theories. The real hope is that in the future, where the power of these machines is expected to increase yet by some orders of magnitude, that it will be practical to utilize these techniques all the time and thereby change the whole style in which much of our theoretical physics is done. For the theorists, the needs are even more insatiable than is the case for the experimentalists. Almost any amount of computing can be utilized to solve these problems in finer detail with relatively little additional effort on the part of the researcher.

Right now there is a computational crisis in the country, documented in the Lax Report. The crisis arises from the lack of access to adequate supercomputer facilities by the university research community. A number of forces are in motion now, attempting to fill the gap. First, there is a national awareness, partly awakened by the Japanese interest in both supercomputers and in artificial intelligence, that there is a need to develop better, more effective supercomputers. Further, some of our theorists have turned into computer architects. In order to solve the lattice gauge problem for high energy physics, a number of university theory groups have now focused their efforts on the building of special-purpose processors to solve their particular problems.

(7) Productivity Enhancement and Other Examples

In terms of our science, there are a number of examples worthy of note with respect to the productivity enhancement properties of these technological developments. The American Physical Society now encourages direct submission of papers in machine-readable form. Although less than 3% of manuscripts are now so submitted, this fraction will certainly increase. For some years now, computer data base indices of physics publications have been developed. The SLAC SPIRES system is widely used, and, with time, as the data base grows in a natural manner it will be more and more effective.

The convenience and efficacy of exchanging information,[6] especially that which is natural to a computer, e.g., programs or data sets from computer to computer, are powerful and will make our efforts much more effective as these features become more standardized. Currently, the major impediments are the lack of standards in operating systems, communications protocols and in the exchange of data and information. We're living in a Tower of Babel as each of us try to take locally developed computer-oriented data and move it to another system at some distant location. That Tower of Babel will be with us for some time, but it is slowly being resolved, and by the end of this decade, some might say sooner, it should be in pretty good shape.

The role that computer technology plays in education is undergoing major changes now. The Computer Aided Instruction (CAI) concepts have been maturing slowly. Many of the original ideas of replacing teachers has finally been recognized by most people to be impractical and undesirable. But, there is a role, a major role, as an adjunct that such approaches are playing and will continue to play in the future. Of course, these functions are in addition to the primary current thrust of computers in the classroom, that of educating and making sure that the current generation of students are computer literate.

With respect to complicated modern experiments, such as collider experiments at CERN and Fermilab, literally dozens of physicists from a number of geographically separated institutions must work together on the same problem. The integration of the computer-based productivity enhancing tools and communications are absolutely crucial for management of such projects and for information transfer amongst these otherwise rather independent individuals trying to work together coherently on a complex enterprise.

These examples are indeed just a sampling. There are many other potential opportunities that may be listed.

E. The Negative Aspects of The Age-The Pitfalls

Certainly there are negative effects, at least from my perception, associated with the age of the computer. There are any number of examples of technical problems or potential problems which one could give here. But I consider most of those minor enough that I will skip over them entirely. Rather it is the sociological problems about which I am most concerned.

I believe a most serious problem that the computing age brings to our research activity is the effect it is having on individuality and in the continuing forcing of nitty gritty technical specializations. The availability of this computing technology makes possible the kinds of very large experiments in which we are forced into bigger groupings to make useful physics contributions. The most blatant example of this is that part of the NASA program which may be called science rather than engineering. Thus, the principle investigators on any given experiment must be involved in organizational enterprises which are certainly by themselves not science, even though they may be a necessary means to do science. Another example from my own field is that the typical leading edge high energy physics experiments today have physicists numbering in the dozens and in a number of cases well over 100. These new large enterprises are typically multi-year and mega-dollar projects, sometimes spanning ten years from the first version of the proposals through the publication of the definitive results from that experiment.

As a corollary to this fact, high energy physics serves as a remarkably good training ground for future professionals in other fields, especially the computing area, and the electronics areas more generally. The numbers of my former students and younger colleagues at Fermilab who are now working for the telephone and computer companies of the world is quite large. These attributes are positive for society and also for physics.

This problem of larger groups and of a specialization in areas supportive of physics rather than directly involved in physics has manifested itself rather recently in theoretical physics also. There are a number of projects throughout the world in which theorists have become quite interested in the detailed inner workings, down to the chip level, of modern electronic equipment. There are strong groups at Cal Tech, Santa Barbara, Columbia and elsewhere who are designing special purpose processors to better solve their current lattice gauge representation. Some of these, although interesting and no doubt capable of making a contribution, will turn out to be larger efforts than the payoffs will have warranted. To make all these things happen, the theorists also have by necessity increased the sizes of their groups.

These are examples of negative facets as I see them. Many may not see these as negative at all. Personally, I long for the good old days when a person could mount an experiment, essentially by himself or with a colleague or two, understand every aspect of it and move it from conception to execution to an extraction of useful physics results in a modest period of time, maybe less than a year.

A general attribute which affects society as a whole is the sociological changes that these group efforts bring with them. The method of communications amongst a number of people, many of whom have their heads glued to a terminal a good fraction of the day, is changing. Rather than face-to-face interactions, there are messages that are sent from one terminal to the other, possibly only across the hall, but often across the country or world. Such communications may be efficient, convenient, typically rather stilted, but whatever they are, they're certainly not personal. I believe that the loss of face-to-face contact is a serious loss or at least a loss that I'm not comfortable with.

A corollary to this is that the sociology of writing a paper is also changing. Now authors type their version of a paper onto a file, edit it, and then pass it on to their colleagues for their editorial changes, additions, comments, etc. Once again, this may be very efficient, but the whole sociology of how a paper is written amongst a number of authors is changing and becoming, to my view, less personal.

The overall direction in which this technology is taking us is more and more to isolate us from our fellow humans and put an intermediary—a digital box, albeit of enormous capability—between us and other members of our society. Every new extension of this technology to make life easier in some dimension imagined as possible, does with effort become possible and the results are a new product or a new tool. There is no control over which tools are developed in our society, nor should there be. Any entrepreneur can choose to invent those tools he believes will be attractive to society. It is from the catalogue of these tools we choose to do our jobs and to make our lives easier. Unfortunately, that brings together a capability which George Orwell feared and so powerfully warned us about in 1948. Now in 1984, although the abridgments of our individualities and freedom that he feared, by most analyses have not come to pass, I believe the potential for a fundamental problem still exists and we should still be very wary of the problem. Further, we are making it possible for some despot who comes to power at some future date to mobilize in very short order this network of powerful tools, which we have made available and use them to less-than-useful ends. Even if that doesn´t happen, I do worry about the depersonalization of the structured society that we are bringing upon ourselves as we advance into this computing age.

F. Conclusion

The computing age may properly be in its middle adolescence period now. Although the growth, which has occurred during the early adolescence period, has been slowing down, with new architectural ideas made possible in part by the technological improvements in the component or semiconductor chip technology, the remarkable growth, in terms of computing power that we have seen during the earlier period, is very likely to continue, probably through the turn of the century. Furthermore, the opportunities and the capabilities that this will bring with it are phenomenal. There is every reason to believe that computations and, in particular, simulations which just a few years ago were unthinkable and had to be done by physical real-world modelling, can and should permeate our industrial society, not just at the research level, but extensively throughout industry, including our smokestack industries. The productivity enhancing tools integrated with powerful computing engines and appropriate software will truly change the way in which we do things in a remarkable way.

At some point, of course, this rapid growth will saturate. Currently, with our silicon-based technology for most of this activity, we're getting close to the limits of optical techniques used in the production of components. Shorter wavelength will help gain another several factors, maybe even an order of magnitude, but the technology will limit somewhere in that vicinity. New materials will take us a bit further, right now gallium arsenide and cooled silicon are the preferred new approaches. Maybe there's yet another order of magnitude or two obtainable from technology.

The new concepts of architecture using multiple processors currently at the level of four or eight CPU's, but eventually rising into the hundreds and possibly eventually the thousands, will be where most of the improvements will occur. The question of how well one can make the large numbers of processors work on the same problem in a general-purpose way is not yet understood. The switching problems are complex, and the software problems have not yet really been adequately faced. Without a doubt, there are some orders of magnitude to be gained with those innovations. In all cases, the final limitation as always will be not the technical, but the economic limit. The computing industry has been growing at a phenomenal rate in terms of its dollar value in our society, but clearly it cannot grow to the level of 100% of our gross national product. That will finally bring us to the natural saturation point. That's not to say that there will not be continued improvements in the performance of machines we build arising from technology. All the other byproducts that come with this technology will, of course, also reap the benefits of these improvements. Before the end of the century, the power of the available tools will be so phenomenal that the opportunities for a whole new class of research will abound. This will certainly change the way we as physicists do our physics and as society more generally does its business.

However, as we pursue our interests and push the frontiers of this computer age technology, we should recall the warning clearly given in George Orwell's "1984" not to allow our individuality to be usurped. I believe his warning is still valid and as we take advantage of the fruits of this fantastic computing age, we should stand above the trees and look at the forest and make sure the forest we're growing is a healthy one.

REFERENCES

1. T. A. Wise, Fortune, p118, September, (1966); T. A. Wise, Fortune, p140, October (1966).

2. There were a few exceptions. In some instances, there were early solid state component failures known as the "purple plague," arising from local high current densities in transistor junctions. These problems were solved relatively quickly by "repopulating" deficient component cards in the affected machines.

3. This was a driving force in Ma Bell´s willingness and desire to undergo divestiture and also for IBM´s major investment in the Satellite Business Corporation some years ago.

4. Private Communication.

5. It is the computer scientists who are pushing hardest in developing the work enhancing tools for their own purposes. This is unlike their lack of real involvement in the other two aspects of computing where they have provided tools for others rather than for themselves and, therefore, have not been involved as users.

6. The airlines are an eloquent example of what can be done merging computers and communications with, in their case, a central data base defining complex schedules in detail down to the last seat, or special meal, etc., with tens of thousands of flights during the course of the year.

TABLE 1

Experimental high energy physics data rates since 1970. Note that at any point in time typical acceptable data rates are defined by the speed which storage media (usually tape) technology will allow. Thus, in 1970, 200 bpi tapes were used, in 1977, 1600 bpi, and in 1984, 6250 bpi tapes are in use. In all cases, not shown, hardwired logic is used to select the triggers of interest to be acquired by the computer and written on tape. In the current scheme, however, the basic rates are so high that a second level of trigger is required. Typically this is tuned to the needs of each experiment, mostly in hardware, but now usually with some software. A third level of trigger processor is now frequently required, typically programmable, requiring tens of microseconds (sometimes as much as hundreds of microseconds) for additional reduction of the raw data.

EXPERIMENTAL HIGH ENERGY PHYSICS DATA RATES

YEAR	RANGE	TYPICAL RATES	
1970			
	10^0-10^3 TRIGGERS/SEC	10^2	10^5 BITS/SEC (TAPE LIMITED)
	10^2-10^4 BITS/TRIGGER	10^3	
1977			
	10^0-10^3 TRIGGERS/SEC	10^2	10^6 BITS/SEC (TAPE LIMITED)
	10^3-10^5 BITS/TRIGGER	10^4	
1984			
	10^3-10^4 TRIGGERS/SEC		
	SPECIAL PROGRAMMED REDUCTION />10		
	PROGRAMMED REDUCTION (50 MIPS) />10		
	10^0-10^2 TRIGGERS/SEC	$2 \cdot 10$	$4 \cdot 10^6$ BITS/SEC (TAPE LIMITED)
	10^4-10^6 BITS/TRIGGER	$2 \cdot 10^5$	

Table 1

FIGURE CAPTIONS

Figure 1: Major computer Execution Bandwidths as a function of time. It is easily seen that for the popular machines produced during the last 30 years the rate of increase of performance capability is slowing down.

Figure 2: A schematic of the data acquisition system for the VLA. Here, it can be seen that there are a number of standard mini-computers, special processors and attached array processors matched to the needs of the radio telescope. Typically, the basic acquisition rate is 300 Kbits/second. The pipeline mapping system is yet another mixture of general-purpose and specialized equipment to perform mapping. The associated DEC 10 off-line system is currently a major bottleneck in the flow of science data.

Figure 3: A schematic drawing of the Princeton Tokamak Fusion Test Reactor (TFTR) control and computer systems. The Cicada system schematically shown is based upon a number of powerful mini-computers configured to operate and control, the Tokamak and to support physics diagnostics for the experimental community.

Figure 4: A schematic drawing of the Livermore double Mirror Fusion Test Facility (MFTF-B) control and diagnostic system. Here the system uses a number of powerful mini-computers for the basic control of the fusion device, including two computers configured for on-line physics analysis.

Figure 5: A schematic drawing of the Space Telescope with its 2.4 meter fused silica mirror. Typical communications rate when in contact with the ground station is 1Mbit/second.

Figure 6: The DUMAND array, shown schematically here, consists of 756 (21x6x6) 16 inch diameter photomultipliers distributed in a 250x250x500M^3 volume 4.7 kilometers deep off the coast of Hawaii.

Figure 7: The electronics configuration for a single plane of DUMAND detectors, consisting of 6 strings of 24 detectors each representing 1/6 of a full detector, is schematically shown here. On each string, 21 Benthos spheres contain a 16 inch photomultiplier, the electronics to drive it and a control microprocessor. The remaining three Benthos spheres are used for another set of experiments. The signals are collected using fiber optic techniques in a Benthos sphere containing the plane processor, which formats the datastream and transmits it to shore 25 kilometers away over 6 fiber optic cables per plane. Each fiber carries continuously a T3 (44Mhz) signal. On shore special-purpose processors backed by standard general-purpose processors collect, decode and analyze the data.

Figure 8: A view of the proton synchrotron at the Fermi National

Accelerator Laboratory. The 2 kilometer diameter main ring is seen in the background and the various experimental areas are in the foreground.

Figure 9: A schematic drawing of the complex control system required to run the Fermilab accelerator. Note a large array of standard mini- and midi-computers as well as a very large number of embedded microprocessors.

Figure 10: The Colliding Detector Facility currently under construction at Fermilab is typical of the current class of new very large detectors being built at accelerator laboratories throughout the world.

Figure 11: An example of an event from the CERN UA1 detector showing the complexity of the events.

Figure 12: A schematic drawing of the Fermilab Colliding Detector Facility data acquisition system. The front-end electronics consists of 150 crates of front-end electronics feeding 150 FASTBUS scanners in the control room. The level 3 processors, which eventually will have a required performance rate of 50 MIPS, will consist of a number of processors with appropriate capability. Each event consists of about 10^5 bits and the recording rate will be limited to between 5 and 10 events per second, the actual rate being tape limited.

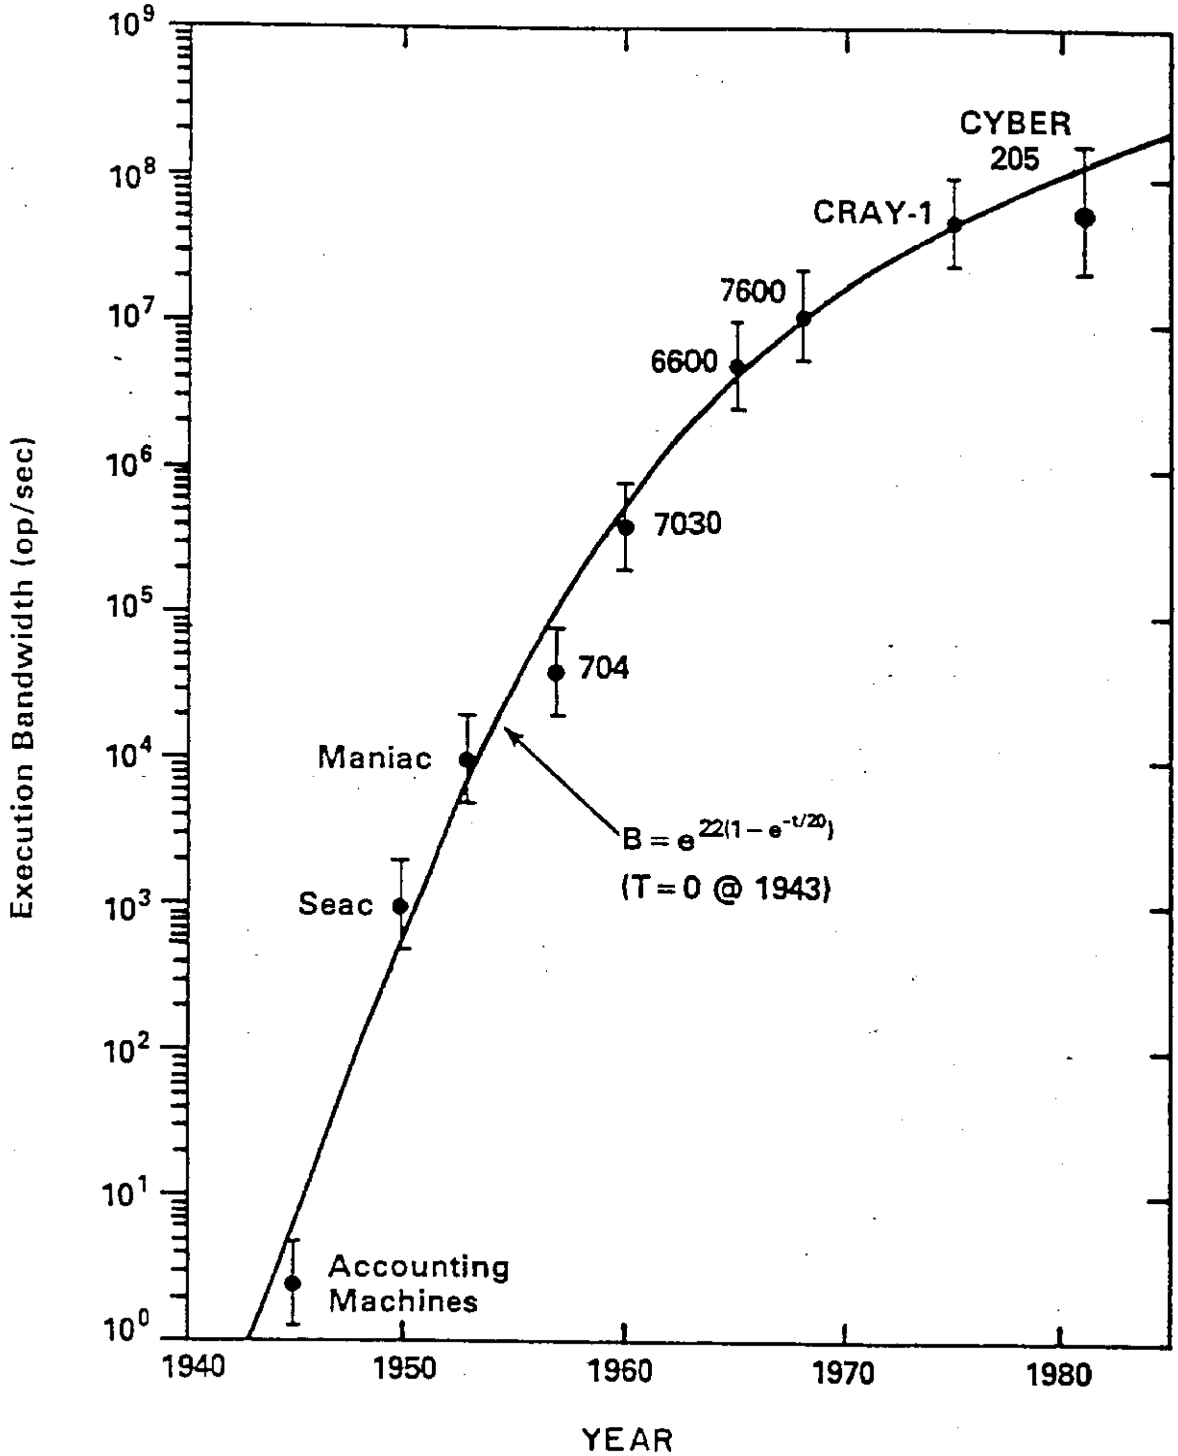

Execution Bandwidth (op/sec)
10^9
10^8
10^7
10^6
10^5
10^4
10^3
10^2
10^1
10^0
CYBER 205
CRAY-1
7600
6600
7030
704
Maniac
Seac
Accounting Machines
$B = e^{22(1 - e^{-t/20})}$
(T = 0 @ 1943)
1940
1950
1960
1970
1980
YEAR

Fig. 1

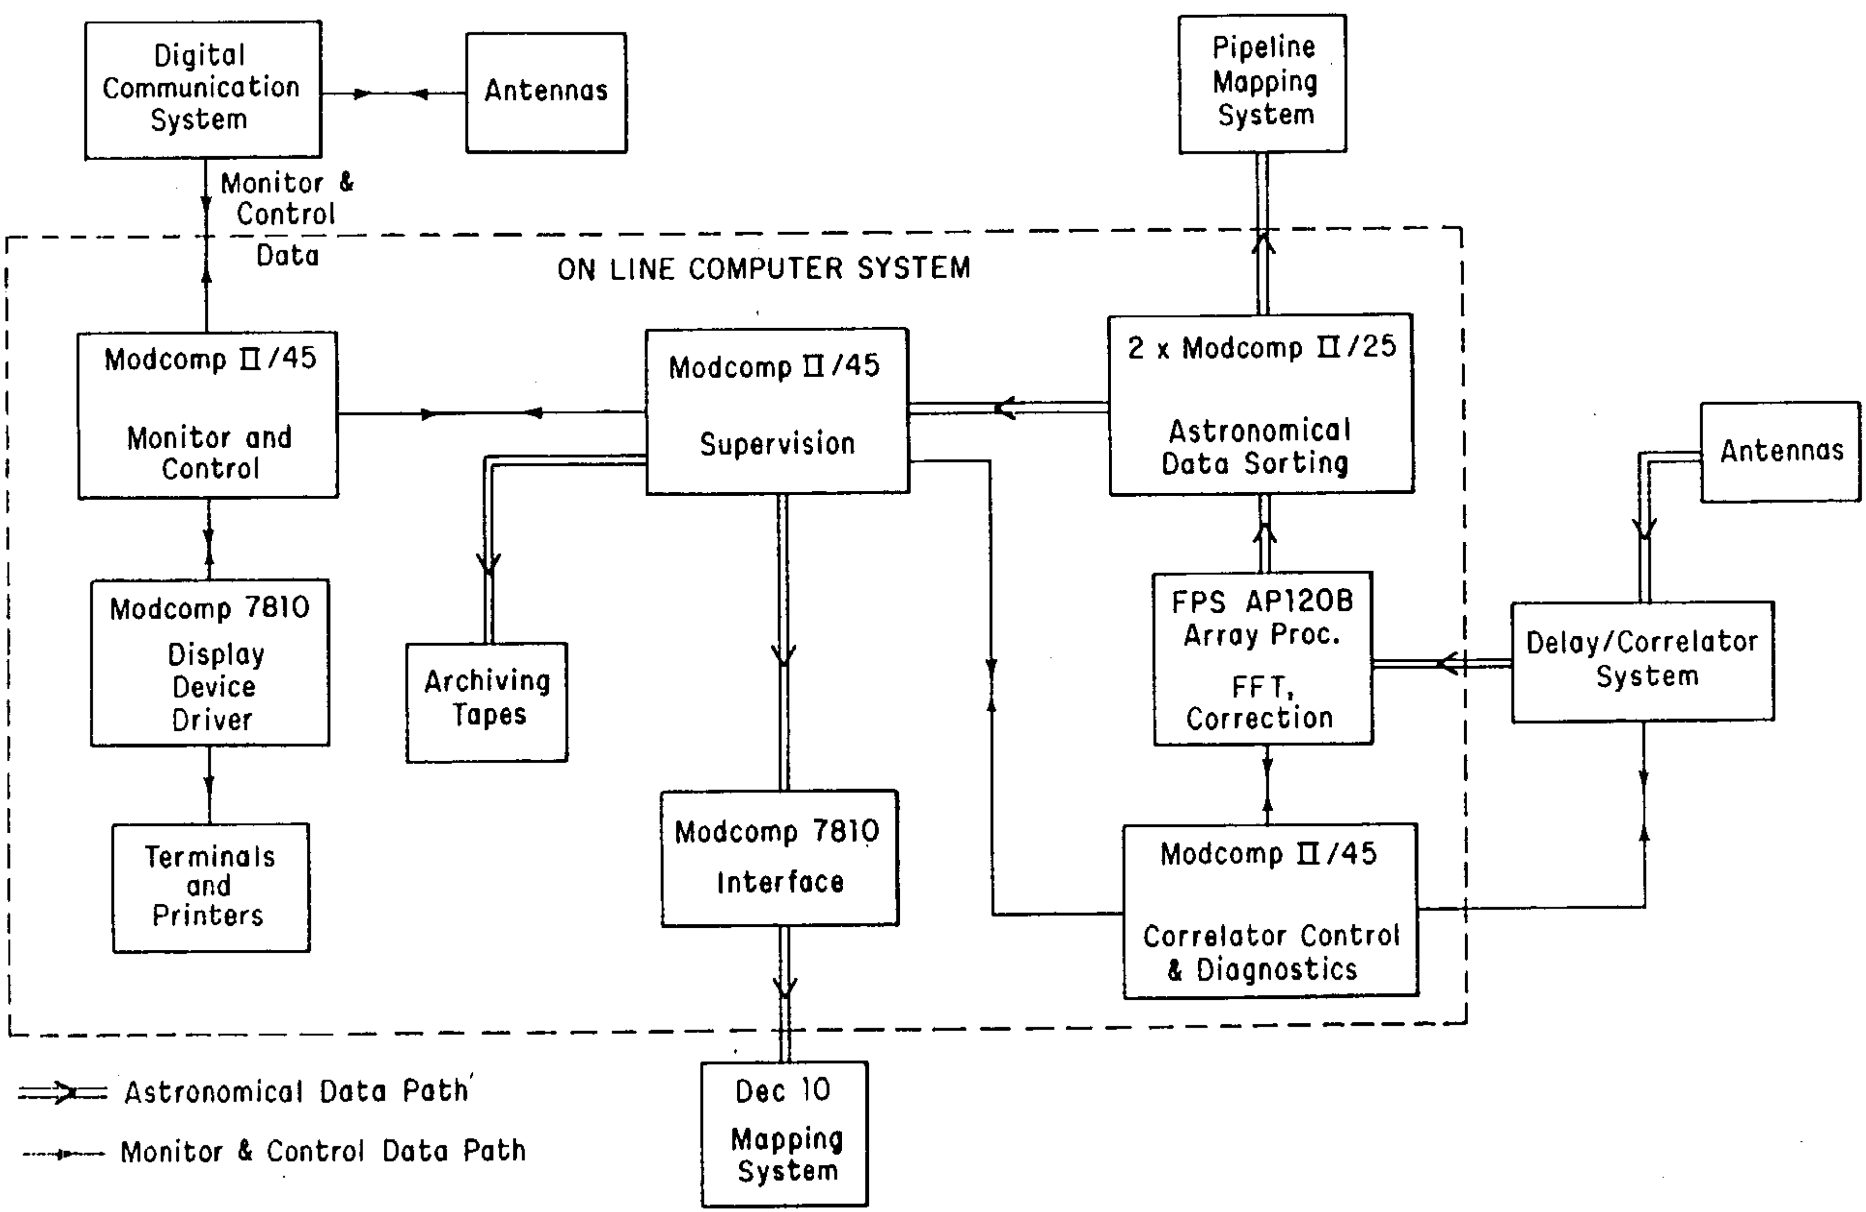

Digital Communication System
Antennas
Monitor & Control Data
ON LINE COMPUTER SYSTEM
Pipeline Mapping System
Modcomp II/45
Monitor and Control
Modcomp II/45
Supervision
2 x Modcomp II/25
Astronomical Data Sorting
Antennas
Modcomp 7810
Display Device Driver
Archiving Tapes
FPS AP120B Array Proc.
FFT, Correction
Delay/Correlator System
Terminals and Printers
Modcomp 7810
Interface
Modcomp II/45
Correlator Control & Diagnostics
Dec 10 Mapping System
Astronomical Data Path
Monitor & Control Data Path

Fig. 2

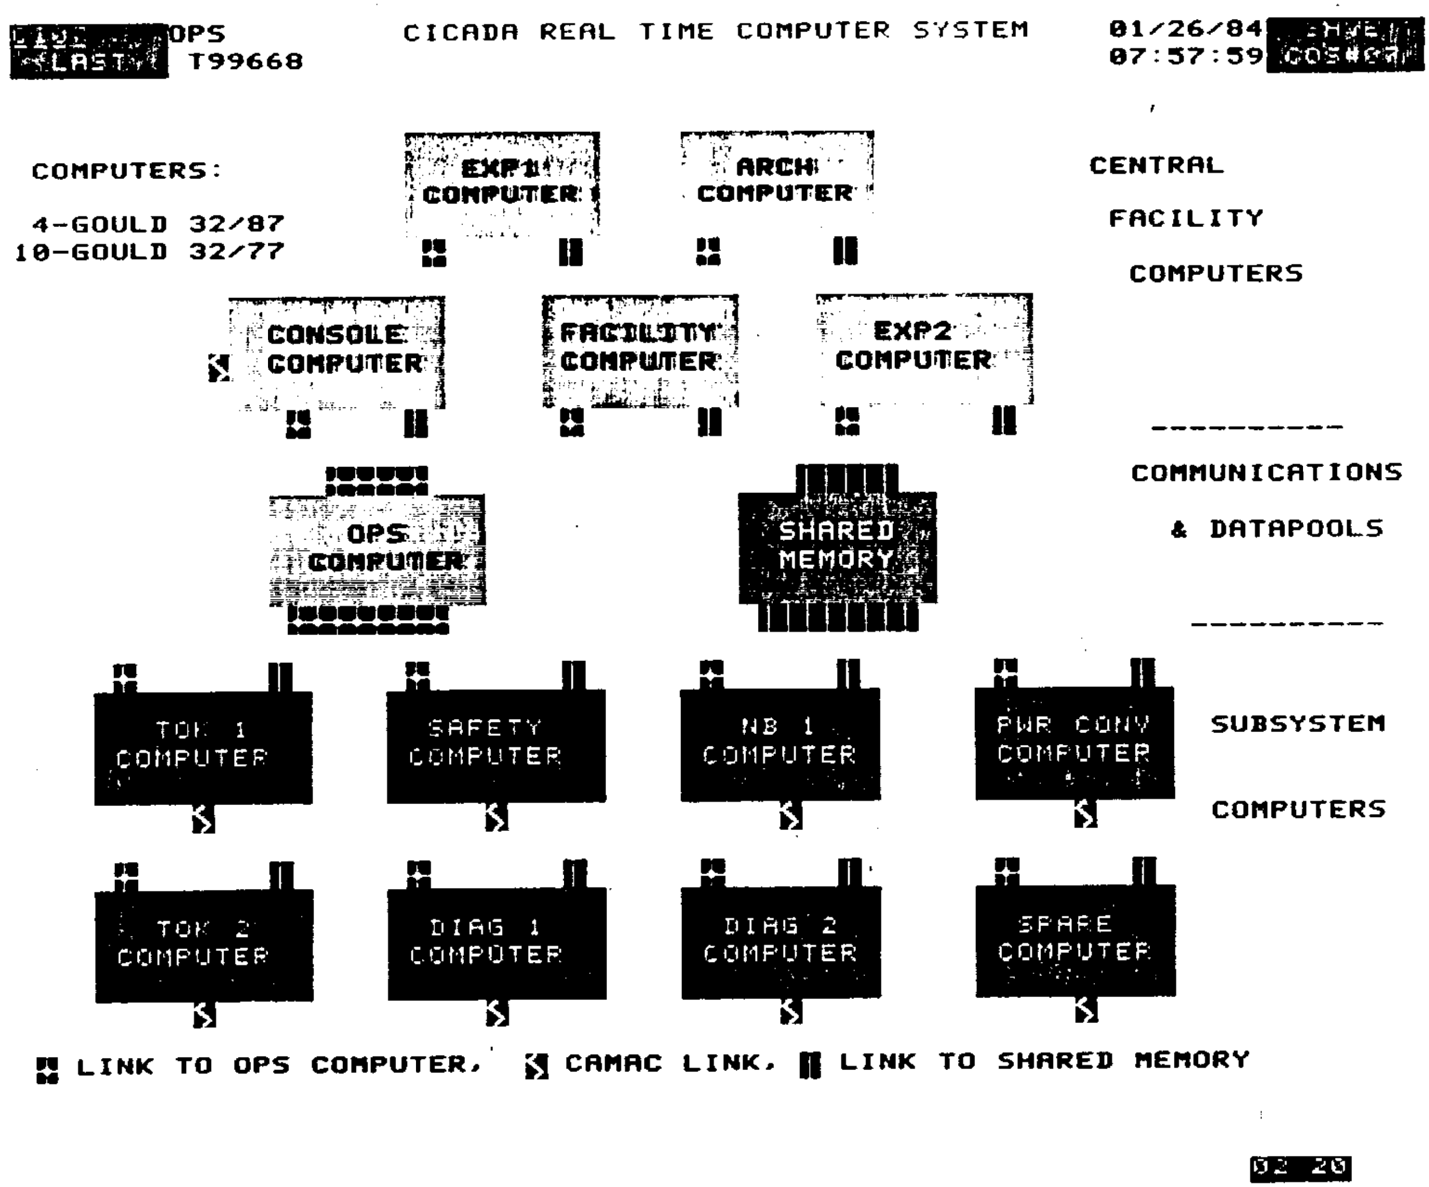
OPS
LAST
T99668
CICADA REAL TIME COMPUTER SYSTEM
01/26/84
07:57:59
COMPUTERS:
4-GOULD 32/87
10-GOULD 32/77
EXP1 COMPUTER
ARCH COMPUTER
CONSOLE COMPUTER
FACILITY COMPUTER
EXP2 COMPUTER
OPS COMPUTER
SHARED MEMORY
TOK 1 COMPUTER
SAFETY COMPUTER
NB 1 COMPUTER
PWR CONV COMPUTER
TOK 2 COMPUTER
DIAG 1 COMPUTER
DIAG 2 COMPUTER
SPARE COMPUTER
CENTRAL FACILITY COMPUTERS
COMMUNICATIONS & DATAPOOLS
SUBSYSTEM COMPUTERS
LINK TO OPS COMPUTER, CAMAC LINK, LINK TO SHARED MEMORY
02 20

Fig. 3

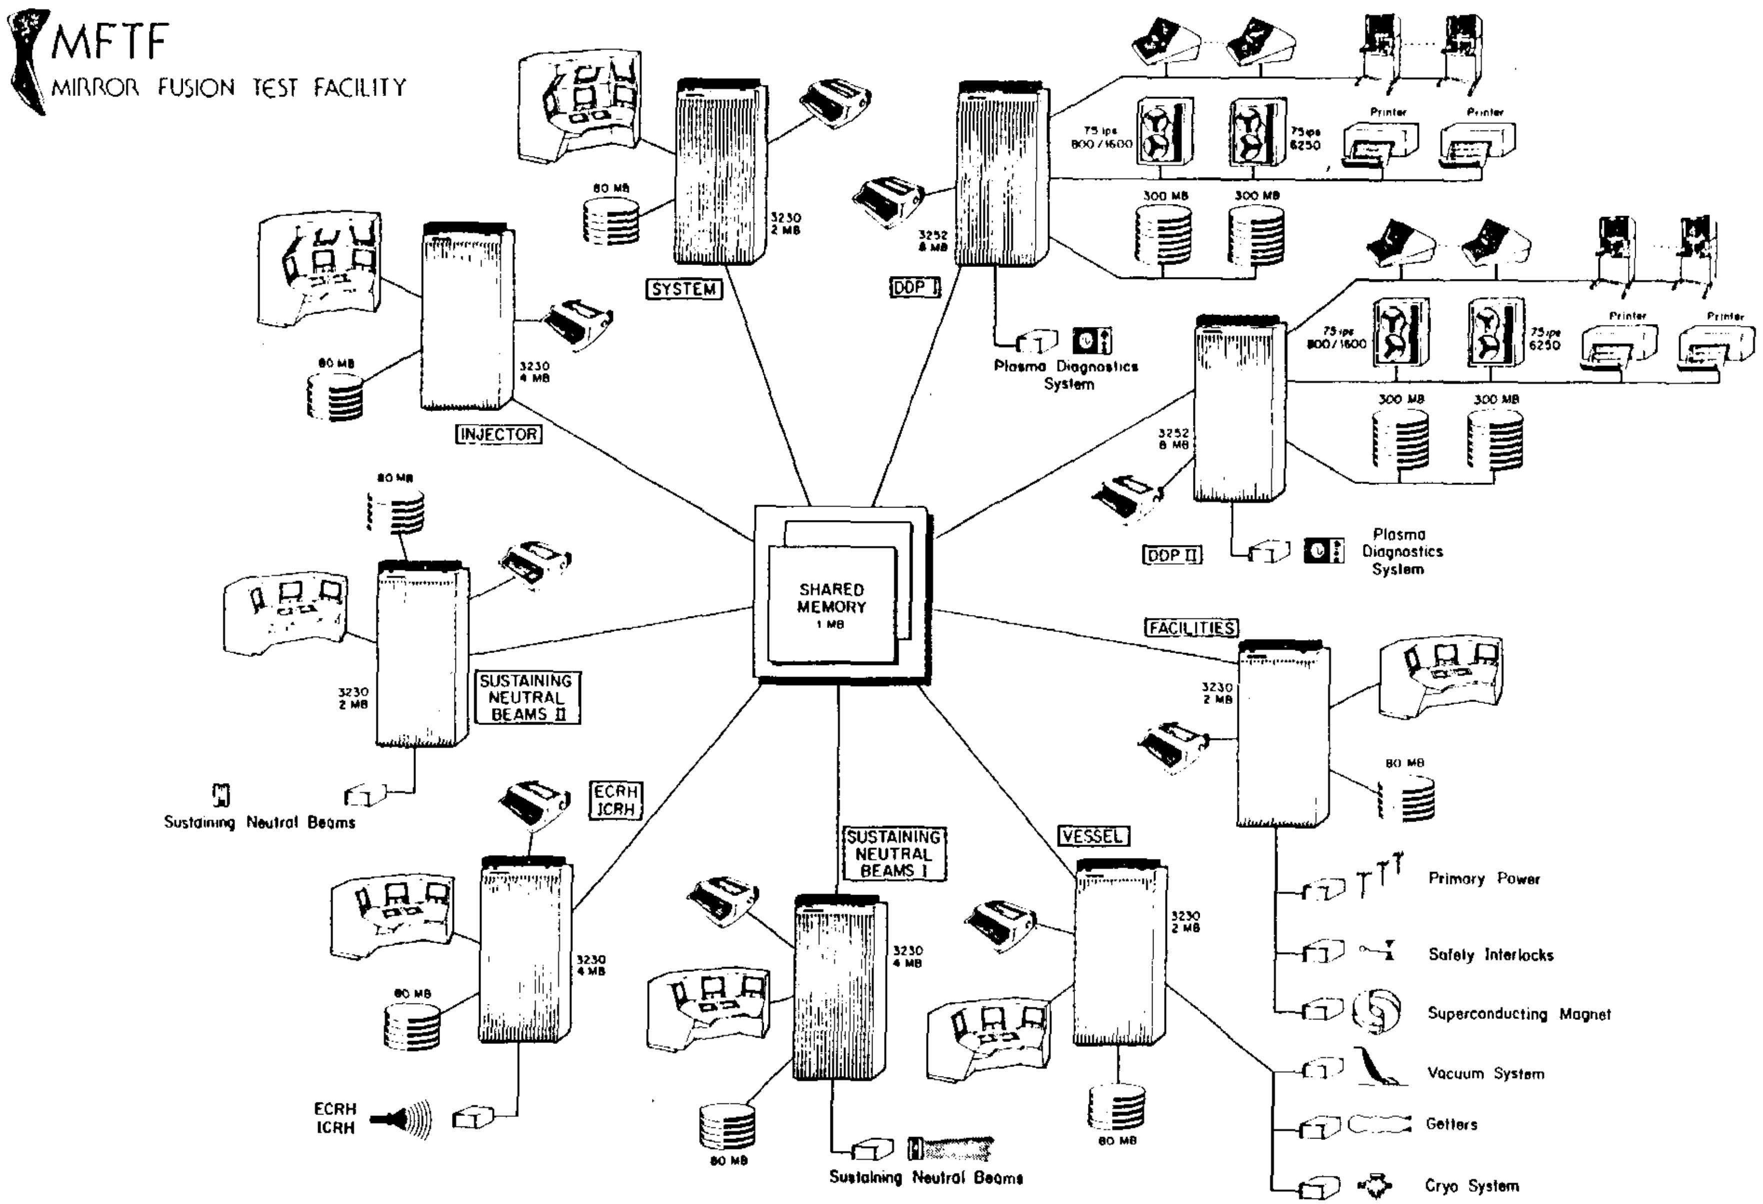

MFTF-B Control and Diagnostics System

Fig. 4

SPACE TELESCOPE CONFIGURATION

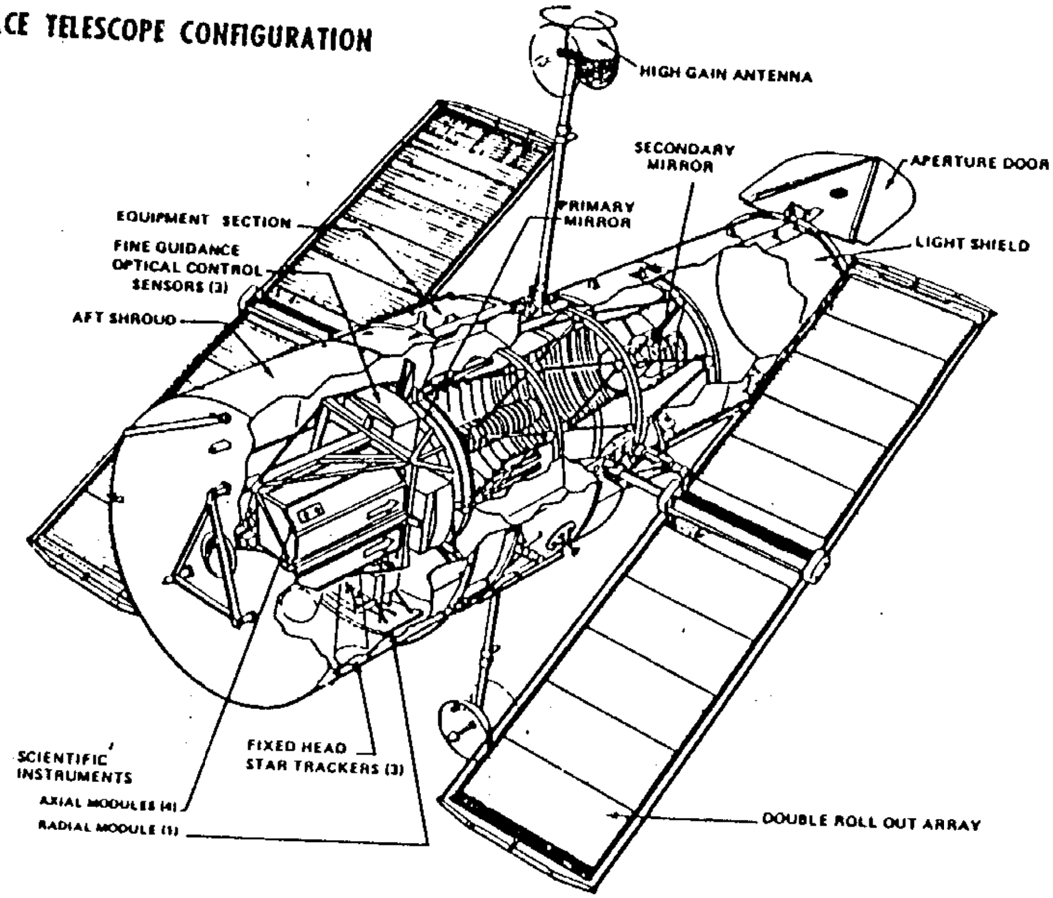

Fig. 5

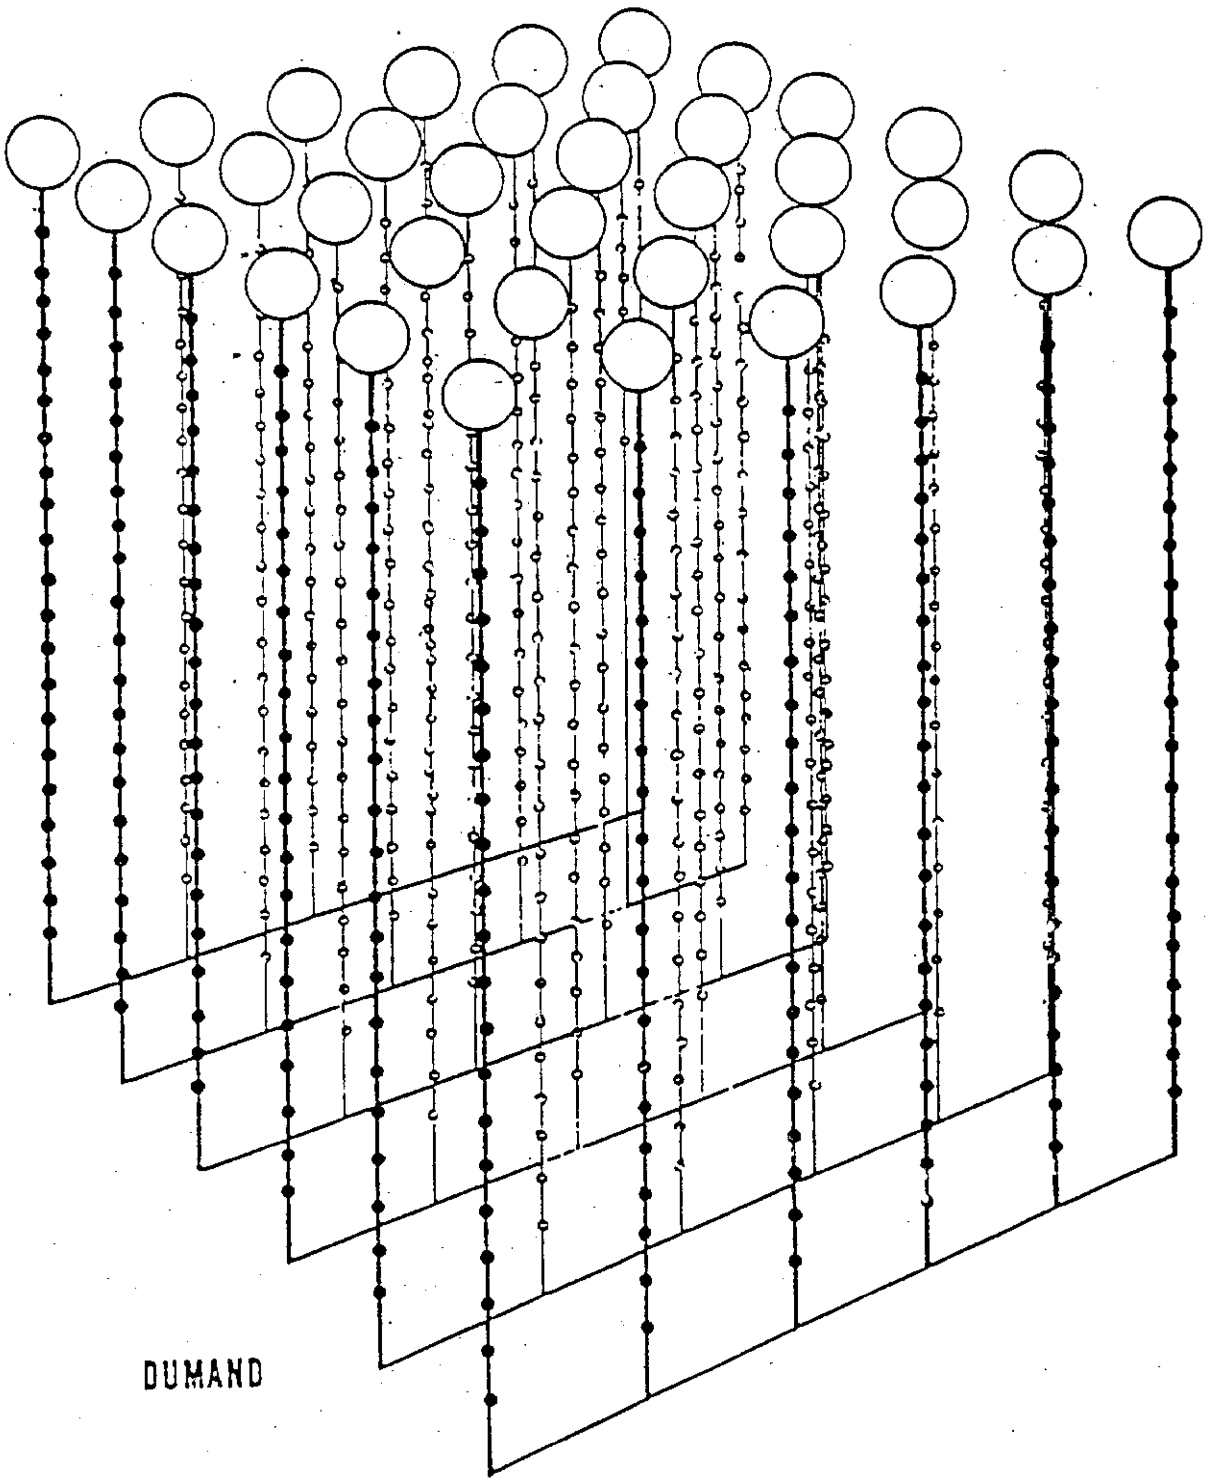

Fig. 6

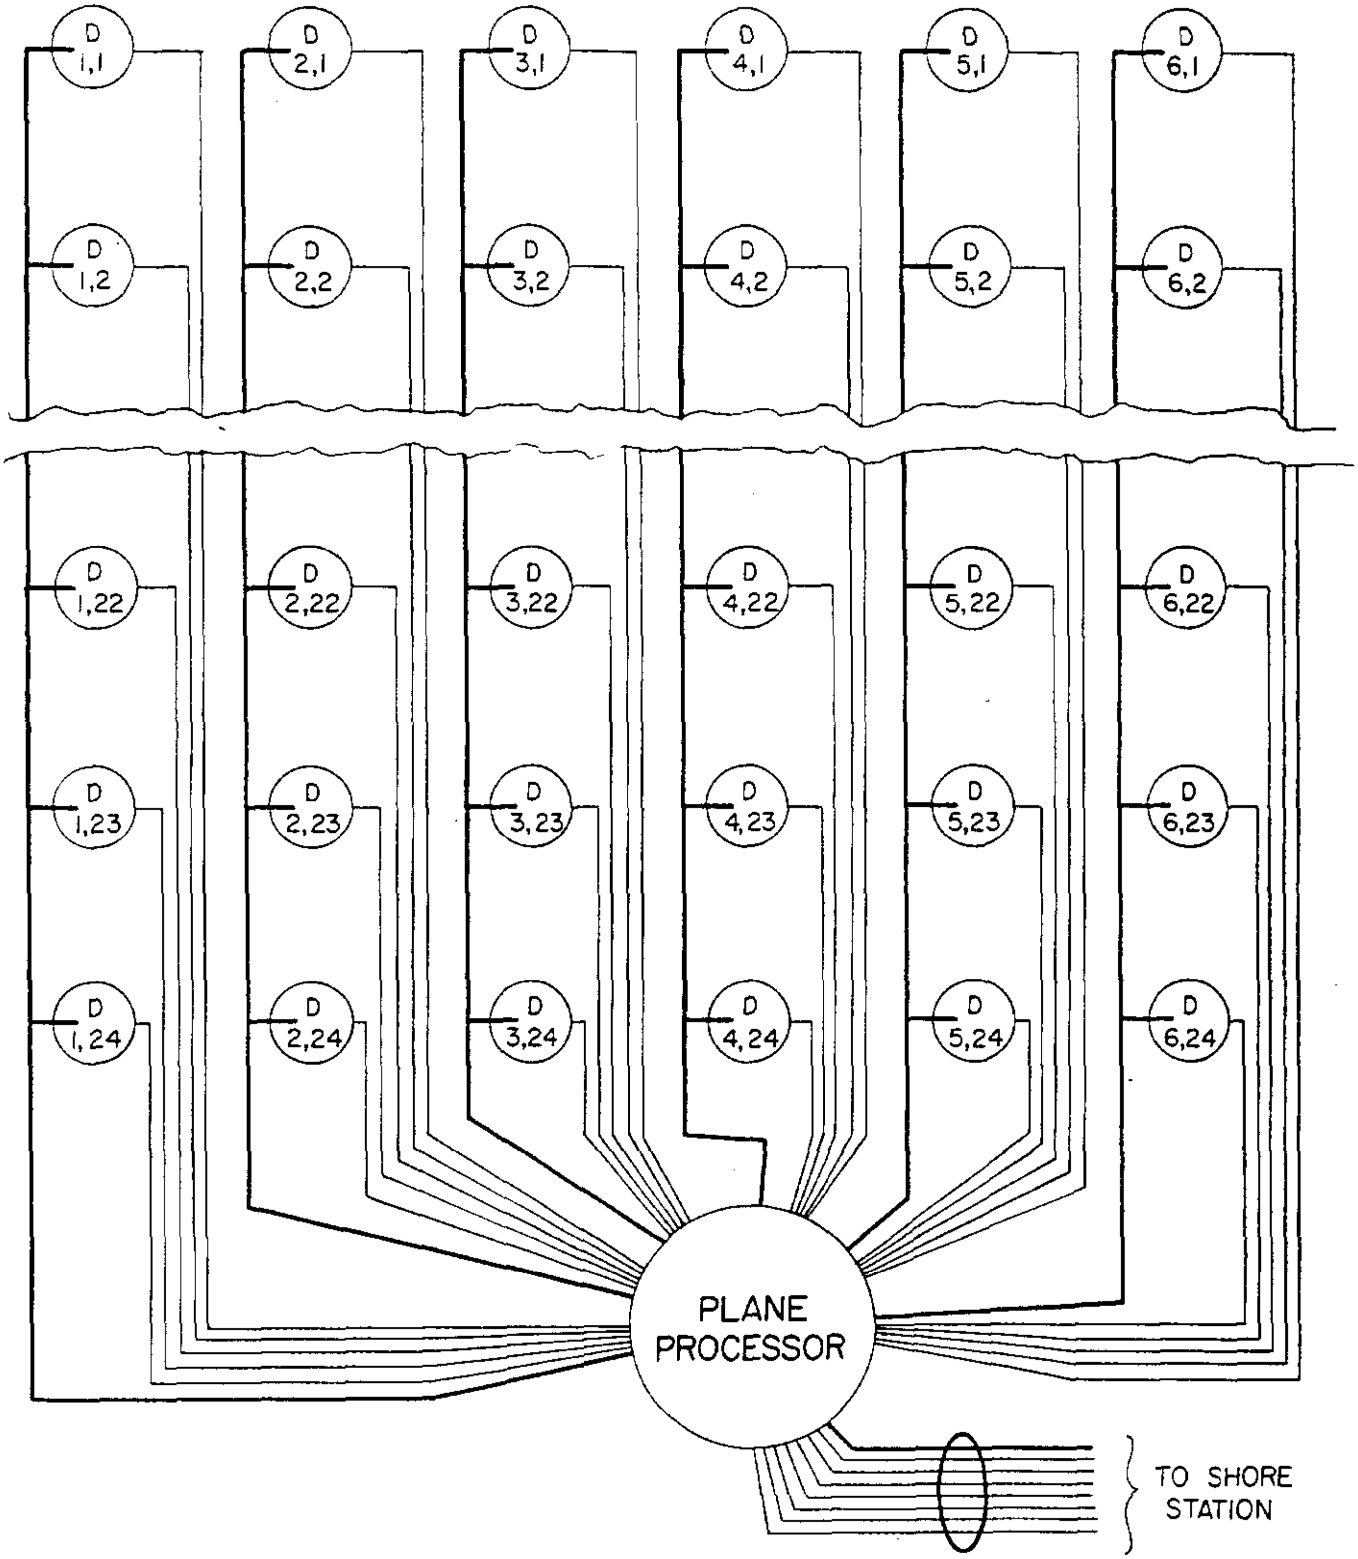
D 1,1
D 2,1
D 3,1
D 4,1
D 5,1
D 6,1
D 1,2
D 2,2
D 3,2
D 4,2
D 5,2
D 6,2
D 1,22
D 2,22
D 3,22
D 4,22
D 5,22
D 6,22
D 1,23
D 2,23
D 3,23
D 4,23
D 5,23
D 6,23
D 1,24
D 2,24
D 3,24
D 4,24
D 5,24
D 6,24
PLANE PROCESSOR
TO SHORE STATION

Fig. 7

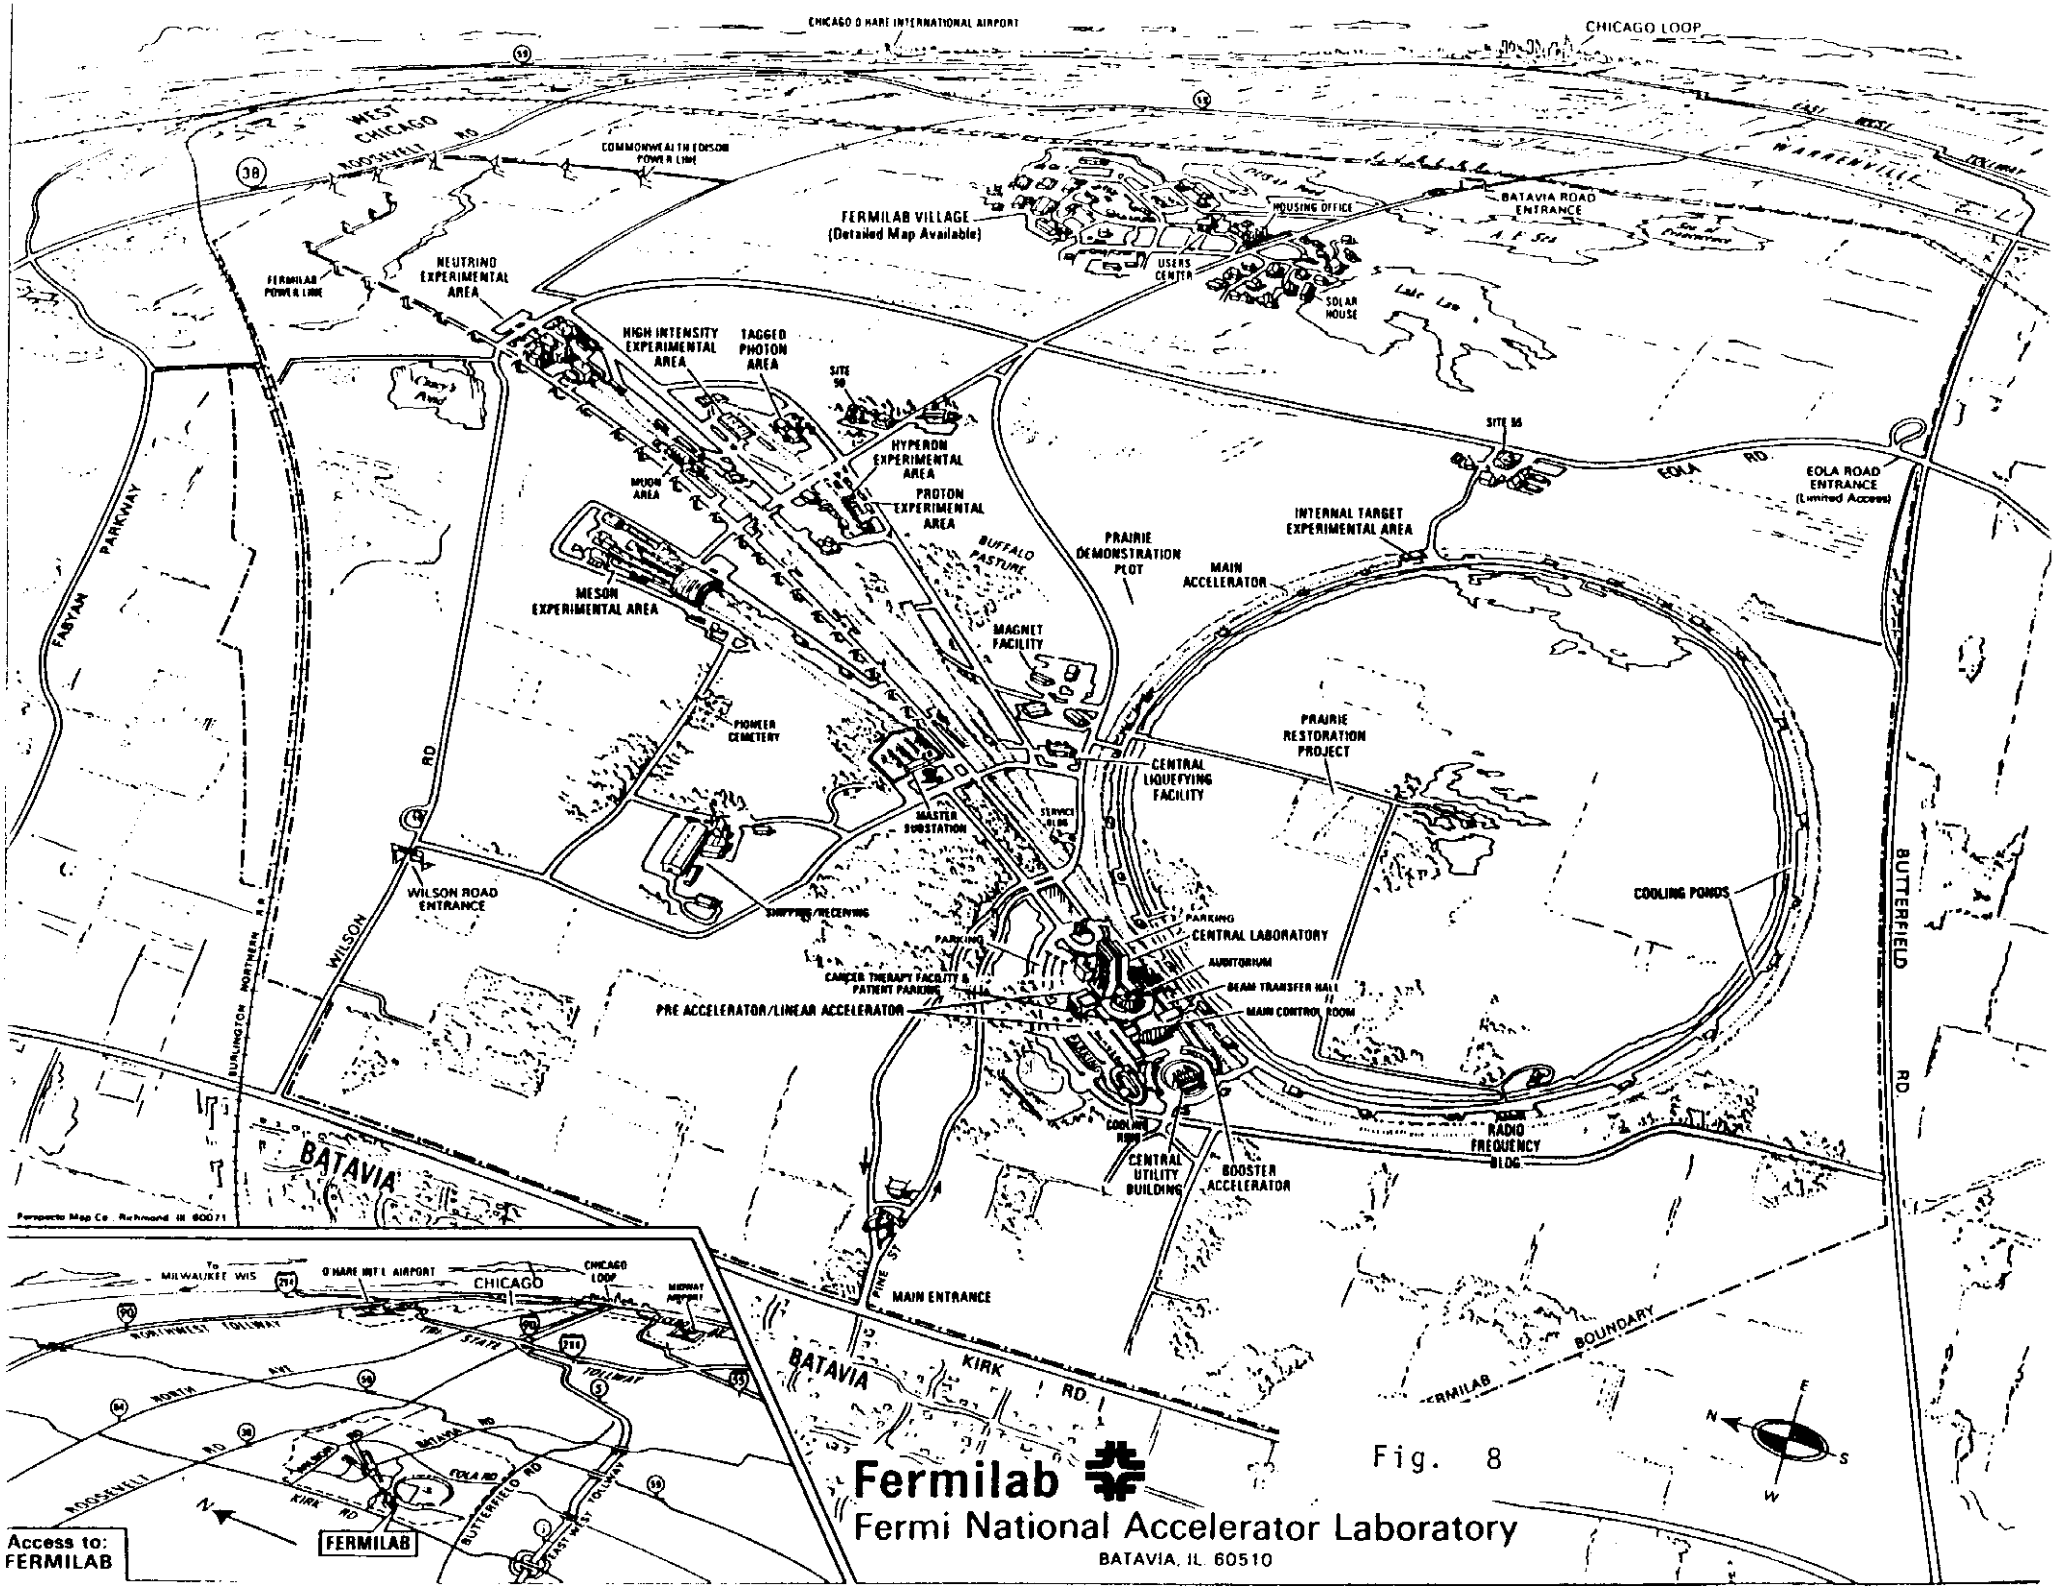
CHICAGO O HARE INTERNATIONAL AIRPORT
CHICAGO LOOP
WEST CHICAGO
ROOSEVELT RD
COMMONWEALTH EDISON POWER LINE
WARRENVILLE
EAST WEST TOLLWAY
FERMILAB VILLAGE (Detailed Map Available)
HOUSING OFFICE
BATAVIA ROAD ENTRANCE
USERS CENTER
SOLAR HOUSE
FERMILAB POWER LINE
NEUTRINO EXPERIMENTAL AREA
HIGH INTENSITY EXPERIMENTAL AREA
TAGGED PHOTON AREA
SITE 50
HYPERON EXPERIMENTAL AREA
PROTON EXPERIMENTAL AREA
MUON AREA
BUFFALO PASTURE
PRAIRIE DEMONSTRATION PLOT
MAIN ACCELERATOR
INTERNAL TARGET EXPERIMENTAL AREA
SITE 55
EOLA RD
EOLA ROAD ENTRANCE (Limited Access)
MESON EXPERIMENTAL AREA
MAGNET FACILITY
FABYAN PARKWAY
PIONEER CEMETERY
PRAIRIE RESTORATION PROJECT
CENTRAL LIQUEFYING FACILITY
SERVICE BLDG
MASTER SUBSTATION
WILSON ROAD ENTRANCE
WILSON RD
SHIPPING/RECEIVING
COOLING PONDS
BUTTERFIELD RD.
PARKING
CENTRAL LABORATORY
AUDITORIUM
BEAM TRANSFER HALL
MAIN CONTROL ROOM
CANCER THERAPY FACILITY & PATIENT PARKING
PRE ACCELERATOR/LINEAR ACCELERATOR
BURLINGTON NORTHERN RR
COOLING RING
CENTRAL UTILITY BUILDING
BOOSTER ACCELERATOR
RADIO FREQUENCY BLDG
BATAVIA
PINE ST
MAIN ENTRANCE
KIRK RD.
FERMILAB BOUNDARY
Perspecto Map Co., Richmond, Ill. 60071
To MILWAUKEE WIS
O HARE INT'L AIRPORT
CHICAGO
CHICAGO LOOP
MIDWAY AIRPORT
NORTHWEST TOLLWAY
TRI STATE TOLLWAY
NORTH AVE
ROOSEVELT RD
BATAVIA RD
EOLA RD
BUTTERFIELD RD
EAST WEST TOLLWAY
KIRK RD
FERMILAB
Access to: FERMILAB
N
E
S
W
Fermilab
Fermi National Accelerator Laboratory
BATAVIA, IL 60510

Fig. 8

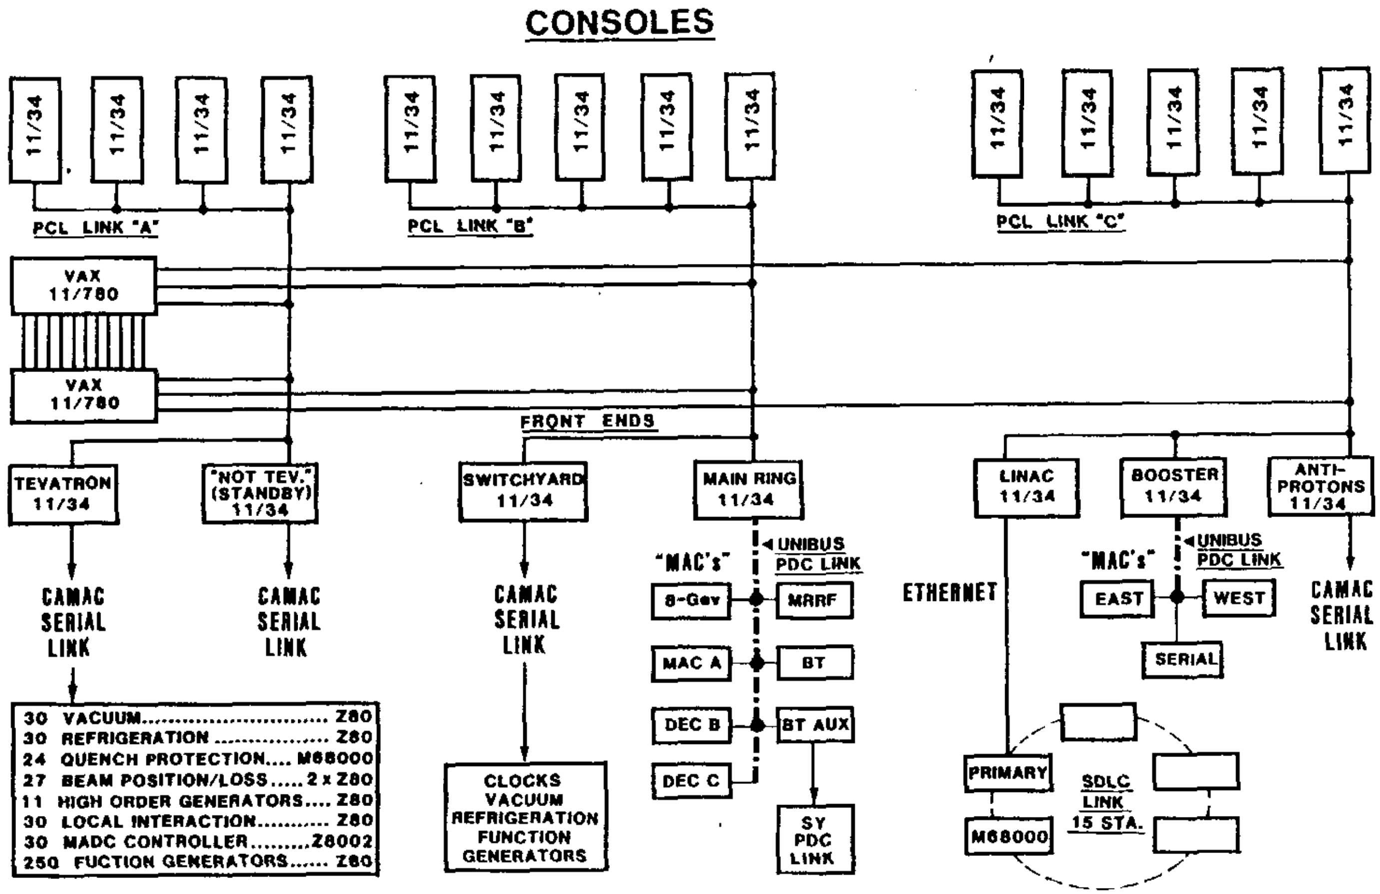
CONSOLES
11/34
PCL LINK "A"
PCL LINK "B"
PCL LINK "C"
VAX 11/780
VAX 11/780
FRONT ENDS
TEVATRON 11/34
"NOT TEV." (STANDBY) 11/34
SWITCHYARD 11/34
MAIN RING 11/34
LINAC 11/34
BOOSTER 11/34
ANTI-PROTONS 11/34
CAMAC SERIAL LINK
CAMAC SERIAL LINK
CAMAC SERIAL LINK
CAMAC SERIAL LINK
30 VACUUM.......................... Z80
30 REFRIGERATION Z80
24 QUENCH PROTECTION.... M68000
27 BEAM POSITION/LOSS..... 2 x Z80
11 HIGH ORDER GENERATORS.... Z80
30 LOCAL INTERACTION........... Z80
30 MADC CONTROLLER.........Z8002
250 FUCTION GENERATORS...... Z80
CLOCKS
VACUUM
REFRIGERATION
FUNCTION
GENERATORS
"MAC's"
UNIBUS PDC LINK
8-Gev
MRRF
MAC A
BT
DEC B
BT AUX
DEC C
SY PDC LINK
ETHERNET
"MAC's"
UNIBUS PDC LINK
EAST
WEST
SERIAL
PRIMARY
M68000
SDLC LINK 15 STA.

Fig. 9

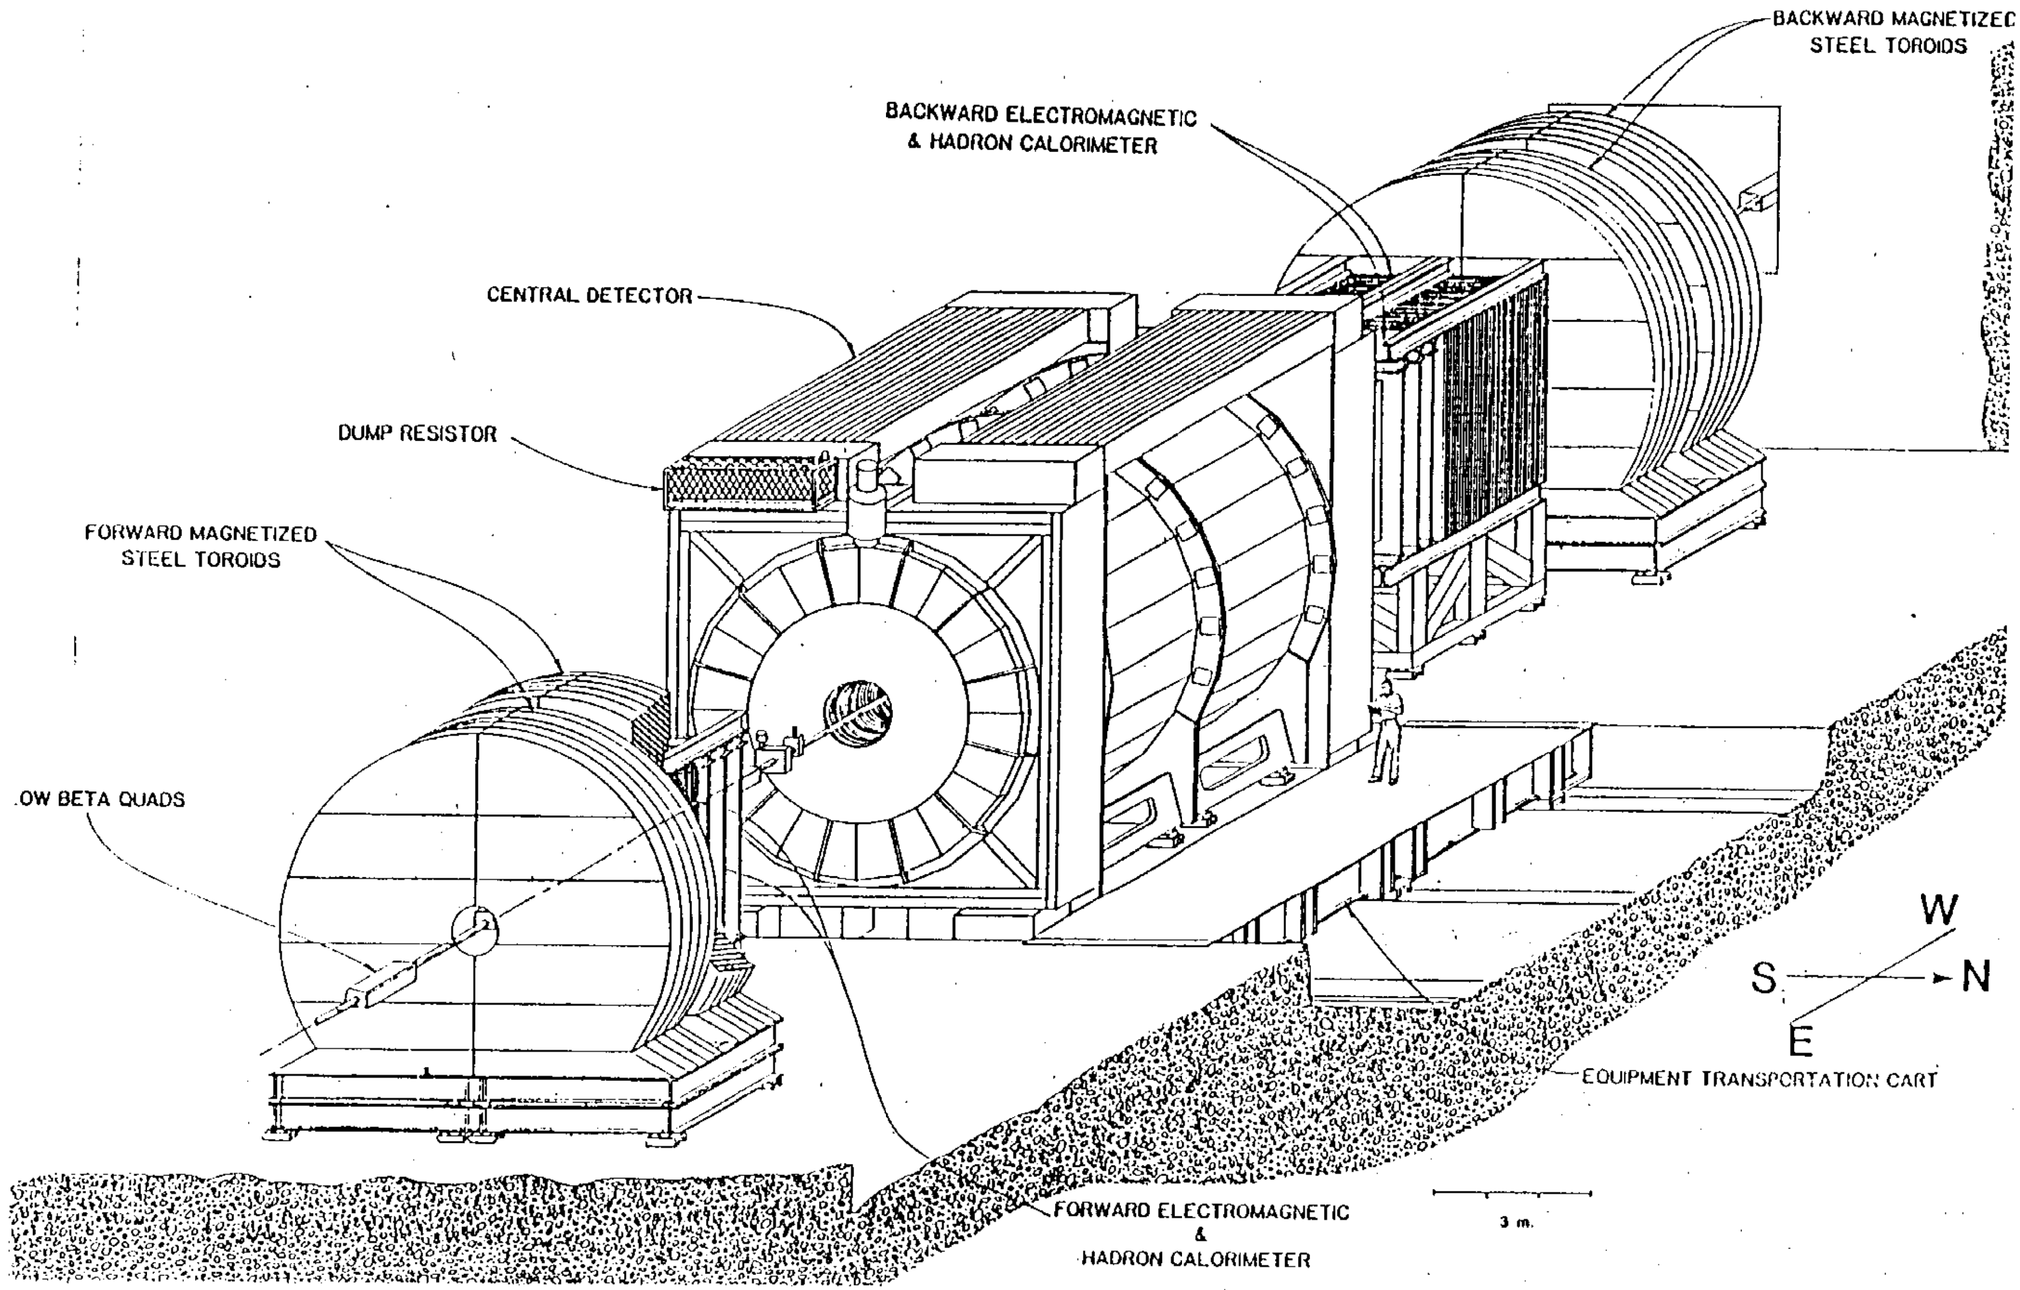
BACKWARD MAGNETIZEC
STEEL TOROIDS
BACKWARD ELECTROMAGNETIC
& HADRON CALORIMETER
CENTRAL DETECTOR
DUMP RESISTOR
FORWARD MAGNETIZED
STEEL TOROIDS
.OW BETA QUADS
W
S
N
E
EQUIPMENT TRANSPORTATION CART
FORWARD ELECTROMAGNETIC
&
HADRON CALORIMETER
3 m.

Fig. 10

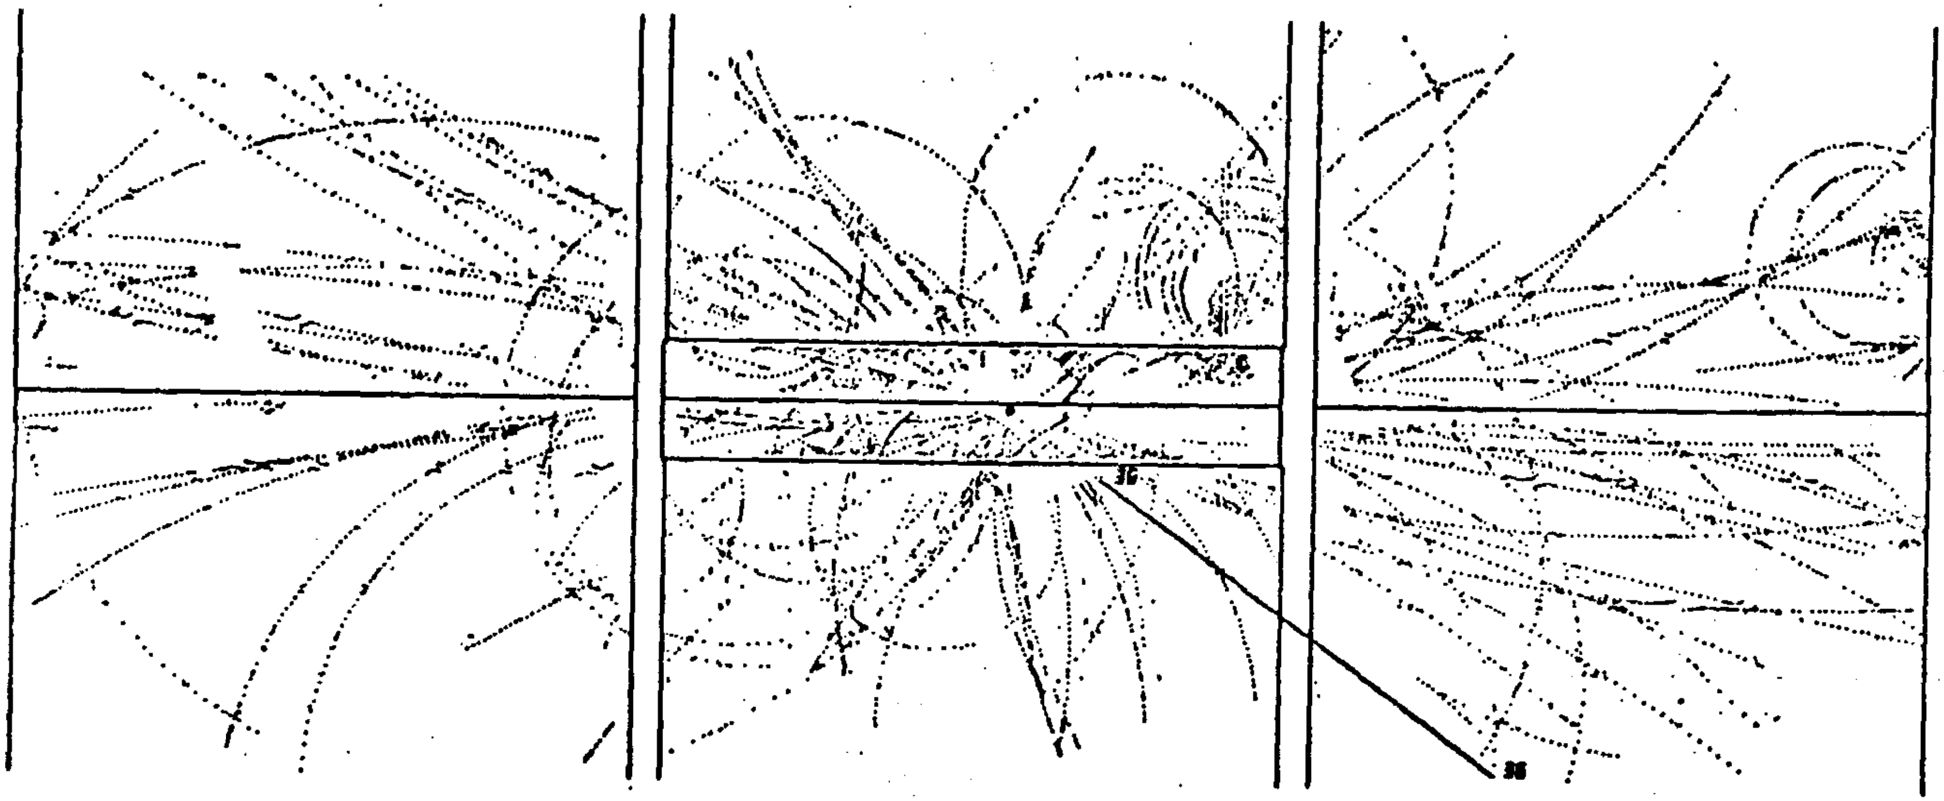

Fig. 11

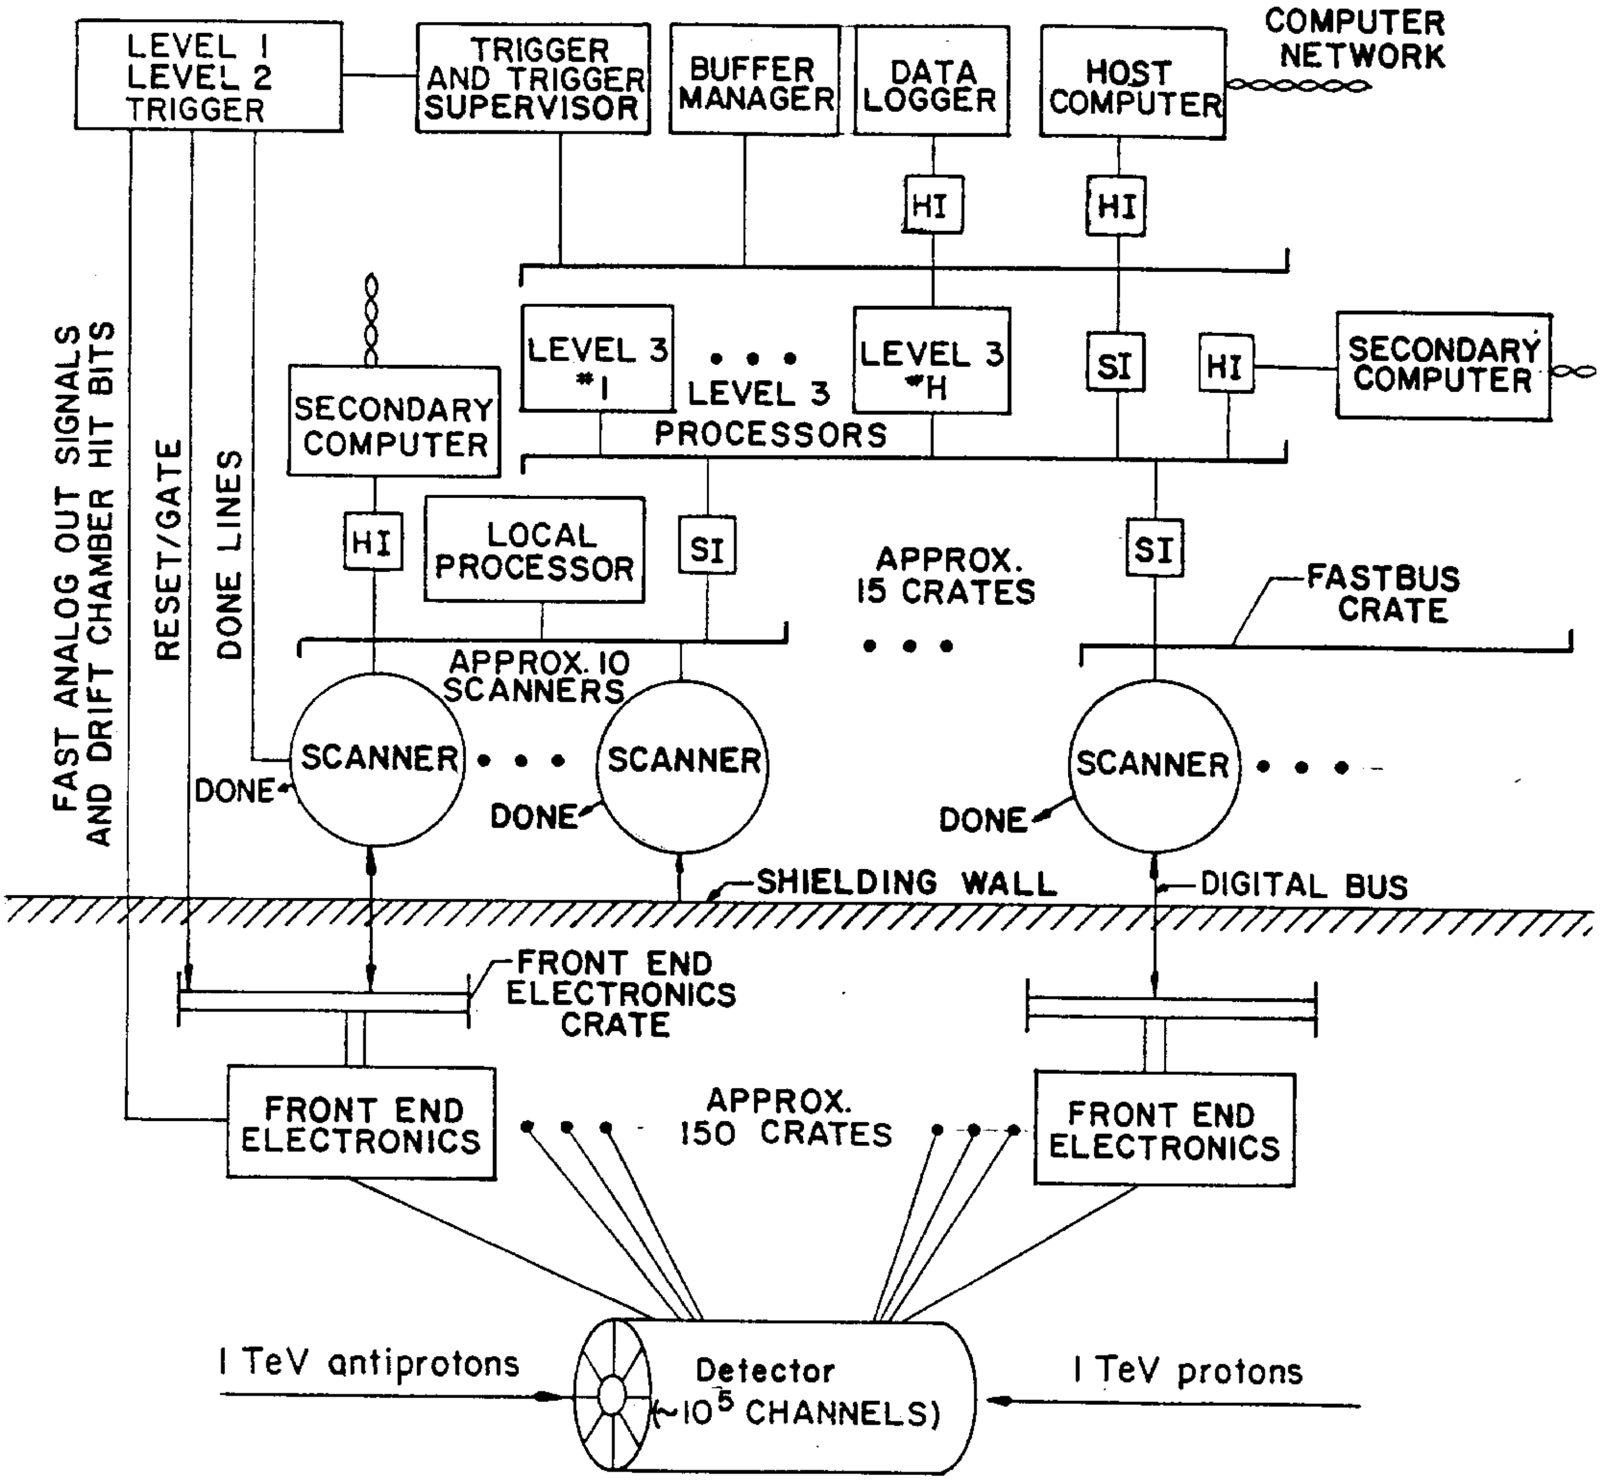
LEVEL 1
LEVEL 2
TRIGGER
TRIGGER
AND TRIGGER
SUPERVISOR
BUFFER
MANAGER
DATA
LOGGER
HOST
COMPUTER
COMPUTER
NETWORK
HI
HI
LEVEL 3
*1
LEVEL 3
PROCESSORS
LEVEL 3
*H
SI
HI
SECONDARY
COMPUTER
SECONDARY
COMPUTER
FAST ANALOG OUT SIGNALS
AND DRIFT CHAMBER HIT BITS
RESET/GATE
DONE LINES
HI
LOCAL
PROCESSOR
SI
APPROX.
15 CRATES
SI
FASTBUS
CRATE
APPROX. 10
SCANNERS
SCANNER
SCANNER
SCANNER
DONE
DONE
DONE
SHIELDING WALL
DIGITAL BUS
FRONT END
ELECTRONICS
CRATE
FRONT END
ELECTRONICS
APPROX.
150 CRATES
FRONT END
ELECTRONICS
1 TeV antiprotons
Detector
(~10^5 CHANNELS)
1 TeV protons

Fig. 12